Intangibles Substitute for Rival Inputs: Evidence from Multinational Expansion*

Sebastian Dyrda
University of Toronto

Guangbin Hong
Michigan State University

Muhammad Ali Sajid
University of Toronto

Joseph B. Steinberg
University of Toronto & NBER

September 14th, 2026

Abstract

Over the last three decades, the intangible intensity of U.S. public firms has nearly doubled and the share of firms producing abroad has more than quadrupled. Non-rivalry ties these two trends together: expanding abroad raises the return to intangible investment, and intangible investment raises the return to expanding abroad. The strength of these channels, and their implications for the macroeconomy, rests on the elasticity of substitution between intangible capital and the rival factors hired locally. We estimate this elasticity on Compustat data using two complementary approaches: (i) a reduced-form approach that instruments multinational activity with the 2004 EU enlargement and bilateral investment treaties; and (ii) a structural approach that uses neural networks to estimate a dynamic heterogeneous-firm model. The reduced-form estimate is 1.8 and the structural estimate is 2.3. Running the instrumental-variables estimation on simulated data from the estimated model reveals that the two estimates are in fact compatible; the reduced-form estimate is a short-run lower bound because intangible capital adjusts slowly. Other experiments with the model show that a reduction in the cost of intangible investment creates spurious measured productivity growth and reduces the labor share.

Keywords: Intangible capital; nonrivalry; elasticity of substitution; multinational production.

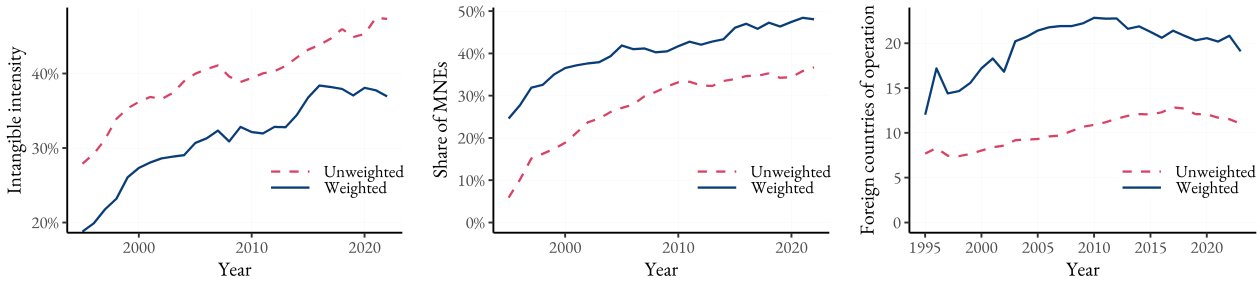
JEL Codes: D24, E22, F23, O33

*Dyrda: sebastian.dyrda@utoronto.ca; Hong: honggua2@msu.edu; Sajid: ali.sajid@mail.utoronto.ca; Steinberg: joseph.steinberg@utoronto.ca. Address: University of Toronto, Max Gluskin House, 150 St. George Street, Toronto, ON M5S 3G7, Canada. We thank Benny Kleinman for his thoughtful discussions. For helpful comments, we thank Chad Jones, Pete Klenow, Chris Tonetti and the participants at the SF Fed Growth Meeting. This research was enabled in part by support provided by Compute Ontario and the Digital Research Alliance of Canada. Computations were performed on the Trillium supercomputer at the SciNet HPC Consortium.

1 Introduction

Intangible capital has become increasingly important in production as firms have expanded their operations across borders. Among U.S. publicly listed firms between 1995 and 2023, intangible intensity—the share of intangible capital in a firm's total capital stock—rose from about a quarter to nearly half, while the share of firms with foreign subsidiaries rose from below 7 percent to almost a third, and the average multinational went from operating in about 7 foreign countries to about 10 (Figure 1). These developments are linked through nonrivalry of intangible capital: the same software, blueprint, or product design can be used in multiple production locations simultaneously. A larger intangible stock raises the profitability of foreign production, while operating in more countries raises the return to intangible investment. How strongly intangible capital and multinational production reinforce each other depends on how firms adjust their input mix as they expand, that is, on the elasticity of substitution between intangible capital and the rival factors that firms employ in each location.

Figure 1: Secular Trends of Intangible Intensity and Multinational Expansion: 1995–2023



(a) Intangible Intensity (b) Share of MNEs (c) Subsidiary Countries

Notes: Data are from Compustat North America and subsidiary disclosures in Exhibit 21 of SEC 10-K filings, 1995–2023.

The central contribution of this paper is to estimate this elasticity. We use two strategies that rely on different sources of variation in the data. The first is a reduced-form instrumental-variables approach that identifies the elasticity from the response of the intangible intensity to policy-driven changes in firm scope. The second is a structural estimation of a dynamic heterogeneous-firm model, where the key moments are the difference in intangible intensity between multinationals and non-multinationals and the correlation between intangible intensity and firm size. Both approaches yield estimates of around two, indicating that intangible capital and rival factors are gross substitutes. The reduced-form estimate is slightly lower than the structural estimate, but we show, by running the reduced-form estimation on simulated data from the model, that adjustment

costs and shocks to the variable cost of multinational production bias the estimate downward in the short run.

We start with a theory in which a firm chooses its scope, its stock of intangible capital, and rival inputs at each location. Each establishment combines the firm's common intangible stock with the rival inputs hired there, and the firm pays an expansion cost of operating abroad that rises with the number of countries it enters. The first-order conditions imply a monotonic relationship between a firm's intangible intensity and its scope, and the slope of this relationship depends only on the elasticity of substitution. The model thus establishes a relationship between intangible intensity and the geographic scope of the firm that can be estimated directly, alongside the relationship between the input mix and relative factor prices on which the literature typically relies to identify the elasticity of substitution.

We take this relationship to the data using Compustat linked to the foreign subsidiaries that firms report in their SEC filings, which provide the number of foreign countries in which each firm operates each year. We start by showing cross-sectional and event-study evidence that intangible intensity rises with scope: multinationals are more intangible intensive than domestic firms in the same industry and of the same size, the gap widens with the number of countries, and intensity rises by about 6 percentage points within two to four years after a firm first becomes multinational, with no increase before entry. We then apply a model-implied instrument to estimate σ: the cost of expanding abroad affects intangible intensity only through scope, so shifts in this cost identify the elasticity. We use two sources of policy variation in this cost, the 2004 EU enlargement, which lowered barriers to foreign direct investment for industries that already had operations in the accession countries, and the spread across host countries over time of bilateral investment treaties. Our preferred specification, which uses both instruments, gives an elasticity of 1.81; the treaty instrument, which has the stronger first stage, gives 1.77 on its own.

Our second approach to estimating the elasticity builds on a quantitative model that extends the theory to a dynamic setting with heterogeneous firms. Relative to the static theory, the model adds persistent and transitory productivity shocks, firm entry and exit, adjustment costs that make firms build their intangible stock gradually, and sunk and per-period costs of foreign affiliates that lead firms to enter and exit foreign production over time. These features let the model speak to both the cross-section and the panel of firms, and account for how firms adjust their intangible stock and their scope to shocks in the short run and in the long run. We estimate the elasticity jointly with the other parameters of the model to match cross-sectional and panel moments from the Compustat data. The structural estimate is 2.25 with a standard error of 0.32.

In both approaches, we can trace the estimate to specific variation in the data. In the IV approach, this is the policy-driven variation in scope described above. In the structural model, the elasticity is pinned down mainly by the intangible-intensity gap between multinational and domestic firms and by the relationship between sales and intangible intensity; neither the IV coefficient nor the event-study responses are among the targets. The gap between multinational and domestic firms reflects both the effect of scope on intensity and the selection of particular firms into multinational production. The model separates the two because other moments, on sales dynamics and on entry into and exit from multinational production, pin down the productivity process and the costs of foreign production, and hence which firms become multinational. We verify that this identification holds over the whole parameter space and not only near the estimate. We estimate the model with the approach of Duarte and Fonseca (2026), which uses a neural network to approximate the mapping from parameters to moments; once trained, the approximation is cheap to evaluate. We can therefore hold the elasticity at any value in its range, re-estimate all other parameters, and record the best attainable fit. The fit deteriorates clearly as the elasticity moves away from 2.25, so within the range we consider, no combination of the other parameters can substitute for it. The two approaches thus identify the elasticity from different variations in the data. The IV estimate uses changes in scope within firms; the structural estimate uses differences in intensity across firms, with the model accounting for which firms expand abroad.

We then use the estimated model to check what the IV design recovers. We draw a panel of firms from the model's stationary distribution and simulate four copies of each: three receive a permanent shock to the costs of foreign production, of increasing size, and the fourth receives none and serves as the control. We estimate the empirical specification on the simulated panel, using a firm's assigned group as the instrument for its change in scope. When the shock is to the fixed cost of operating abroad, the estimate is about 1.95 one year after the shock and about 2.3 twenty years later. The design recovers the elasticity once firms have adjusted their intangible stocks. When the shock is to the variable cost of foreign production, the estimate is 1.6 to 1.7 in the short run and 2.9 to 3.0 in the long run. A change in the variable cost also changes how much labor and tangible capital affiliates hire, which changes the return to intangible capital even for firms whose scope does not change, so the instrument affects intensity through a channel other than scope. We do not know whether the EU enlargement and the investment treaties changed fixed costs, variable costs, or both. In all cases, however, the short-run estimate lies below the long-run elasticity, and our empirical estimates fall within the short-run range. We therefore read them as a lower bound on the long-run elasticity.

The elasticity also matters for two macroeconomic questions. One is how much of measured productivity growth reflects the accumulation of intangible capital and how much reflects improvements in efficiency (Corrado, Hulten, and Sichel, 2009; Crouzet and Eberly, 2021). The other is why the labor share has fallen as investment goods have become cheaper (Karabarbounis and Neiman, 2014) and production has shifted toward firms with low labor shares (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020). We examine both with an experiment in the estimated model: an unanticipated, permanent 10 percent reduction in the price of intangible investment, with wages and rental rates held fixed. We compare the estimated technology with a Cobb–Douglas benchmark. Consider productivity measurement first. When the elasticity is above one, the output elasticity of intangible capital rises with a plant's intangible intensity. Growth accounting with common output elasticities therefore attributes part of the growth in intangible capital to productivity, even when intangible capital is measured and included as an input. In the experiment, this adds about 0.5 percentage points to measured productivity growth over five years, although productivity itself does not change. Under Cobb–Douglas the error is zero. Consider the labor share next. Because intangible capital substitutes for tangible capital and labor, more intangible capital lowers the share of a plant's output paid to labor. At our estimate it still raises output and employment, so the labor share falls because output grows faster than payroll. In the experiment, domestic output rises by 2.6 percent under the estimated technology and by 2.4 percent under Cobb–Douglas. Domestic employment rises by 1.2 percent under the estimated technology and by 2.4 percent under Cobb–Douglas. The labor share falls by 0.8 percentage points, almost entirely through changes within plants and not through reallocation across them. Under Cobb–Douglas it does not change. Cheaper intangibles raise output by about the same amount under both technologies. The difference is in how firms produce it: under the estimated technology, part of the increase comes from intangible capital instead of labor.

Related Literature. Our paper contributes to the literature on multinational production. Since Helpman, Melitz, and Yeaple (2004), this literature models foreign production as the choice of the most productive firms, and quantitative frameworks find large gains from the cross-border transfer of firm-specific technologies (Ramondo and Rodríguez-Clare, 2013; Tintelnot, 2017).[1] Alviarez, Cravino, and Ramondo (2023) isolate the firm-embedded component of productivity, such as blueprints, management practices, and intangible capital, that multinationals transfer across borders, and find that it accounts for roughly one-third of cross-country income differences. In these frameworks the transferable input is a productivity draw that the firm takes as given. In the

[1] Ramondo, Rodríguez-Clare, and Tintelnot (2015) document the stylized facts of multinational production; Antràs and Yeaple (2014) survey the literature.

technology-capital models of McGrattan and Prescott (2009, 2010), firms accumulate the non-rival input, but it enters production in Cobb–Douglas form, so the number of locations does not affect the firm's input mix. We endogenize the accumulation of the non-rival input and estimate its substitutability with local inputs, so that a firm's intangible stock and its scope are jointly determined. A related set of papers studies innovation within multinational firms. Arkolakis, Ramondo, Rodríguez-Clare, and Yeaple (2018) model the separation of innovation from production across countries, and Bilir and Morales (2020) estimate that the median U.S. multinational realizes abroad one fifth of the return to its domestic R&D, direct evidence that intangible investment at headquarters raises output at foreign affiliates. Our focus is on how this non-rival input combines with local inputs once it exists.

We also build on studies of intangible capital. Corrado et al. (2009) introduce the aggregate measurement of intangible investment, and Peters and Taylor (2017) develop the firm-level measure of intangible stocks from capitalized R&D that we adopt.[2] Crouzet, Eberly, Eisfeldt, and Papanikolaou (2022) identify non-rivalry and scalability as the properties that distinguish intangibles from other capital, Jones and Tonetti (2020) formalize non-rivalry as a source of increasing returns, and Crouzet, Eberly, Eisfeldt, and Papanikolaou (2026) study how the scope over which a firm deploys its intangibles affects growth and the measurement of capital services.[3] In these studies the substitutability between intangibles and rival inputs is an assumption. We estimate it, and we show that the geographic scope of production identifies it.

Finally, our paper is related to estimates of the elasticity of substitution between intangible inputs and other factors. Aum and Shin (2024) and Jerbashian (2024) estimate elasticities between software or intellectual-property capital and labor above one, while Lashkari, Bauer, and Boussard (2024) find an elasticity below one for IT capital.[4] These studies identify substitution from variation in factor prices. The closest paper to ours is Kleinman (2025), who uses the geographic expansion of multi-establishment firms within the United States to estimate an elasticity of 1.3 between headquarter services, a non-rival input, and local rival factors. We use the same source of

[2] Eisfeldt and Papanikolaou (2013) measure organizational capital from SG&A expenditures; O'Mahony, Corrado, Haskel, Iommi, Jona-Lasinio, and Mas (2018) extend aggregate intangible measurement to European economies; and Crouzet and Eberly (2021) link intangible mismeasurement to productivity and markups.

[3] A related literature connects the scalability of non-rival inputs to rising concentration, declining dynamism, and technology diffusion through market expansion (Perla, Tonetti, and Waugh, 2021; Benhabib, Perla, and Tonetti, 2021; Aghion, Bergeaud, Boppart, Klenow, and Li, 2023; Klenow and Li, 2025).

[4] Aum and Shin (2024) note that the IT capital measured by Lashkari et al. (2024) combines non-rival inputs (e.g., software) with rival ones (e.g., computers). For the elasticity of substitution between tangible capital and labor, see Karabarbounis and Neiman (2014), who estimate an elasticity above one from the declining relative price of investment goods, and Oberfield and Raval (2021), who aggregate plant-level substitution to an elasticity below one.

identification across borders, where the number of countries in which a firm produces is observed directly and discrete policy changes to FDI barriers provide instruments for it. We also estimate a dynamic model of the firm that delivers a structural counterpart to the IV estimate, shows that the IV estimate understates the long-run elasticity when intangible capital adjusts slowly, and lets us study what the elasticity implies for productivity measurement and the labor share.

2 Theory

We develop a theory in which firms jointly choose how much intangible capital to accumulate at their headquarters and the geographic scope of their production. The theory shows how the non-rivalry of intangible capital links a firm's global scope to its intangible intensity, and it delivers a closed-form relationship between the two whose slope identifies the elasticity of substitution. We take this relationship directly to firm-level data in Section 3. The notation follows the quantitative model of Section 4, of which the theory is a static counterpart.

2.1 Environment

There is one home country and N identical foreign countries. In the home country there exists a set of firms that differ in productivity $a \in \mathcal{A}$. We index each firm by its productivity, as all firms with the same productivity make identical decisions. For simplicity, we assume there is no international trade. To produce in foreign countries, firms set up affiliates abroad, subject to a multinational expansion cost that depends on the number of foreign countries served, $\kappa(n)$. Firms also invest in intangible capital z at price P. Intangible capital is non-rival and can be used in all locations. Each location produces a final good using the local rival input ℓ and the firm-level non-rival input z. The single rival input ℓ stands in for the bundle of labor and tangible capital, $g(k, \ell)$, of the quantitative model; collapsing the bundle to one input simplifies the derivations without affecting the mechanism. All firms take the price of the rival input W, the price of intangible capital P, and the output price p as given.

A firm with productivity a, intangible capital z, and n foreign affiliates earns global static profits

$$\pi(a, z, n) = \max_{\ell} \{(1+n)[p f(a, z, \ell) - W\ell]\} = (1+n)\tilde{\pi}(a, z), \tag{1}$$

where $\tilde{\pi}(a, z) \equiv \max_{\ell} \{p f(a, z, \ell) - W\ell\}$ is the profit earned in a single location and the same rival input is chosen in every location because countries are identical. The value of a firm with

productivity a is

$$V(a) = \max_{z,n} \{\pi(a, z, n) - Pz - \kappa(n)\}. \tag{2}$$

We define intangible intensity as the expenditure share on intangible capital, $s^z \equiv Pz/[(1+n)W\ell + Pz]$. We make the following assumptions on the production technology and the expansion cost.

Assumption 1. *Production and cost functions are as follows:*

1. *The production technology features constant elasticity of substitution between the two factors, that is*

$$f(a, z, \ell) = a\left[\mu^{1/\sigma} z^{\frac{\sigma-1}{\sigma}} + (1-\mu)^{1/\sigma} \ell^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\theta\sigma}{\sigma-1}}$$

   where θ is the returns-to-scale parameter, σ is the elasticity of substitution between z and ℓ, and μ controls the importance of intangible capital in the CES production function. We assume $\theta < 1$.

2. *The expansion cost is iso-elastic in the number of foreign countries, that is*

$$\kappa(n) = \bar{\kappa}\, n^{\gamma}$$

   where $\bar{\kappa} > 0$ governs the level of the cost function, and $\gamma > 1$ is the cost elasticity with respect to the number of foreign affiliates. We assume the marginal cost of multinational expansion increases faster than global profits as the firm operates in more countries, i.e. $\gamma > \frac{\partial \ln \pi(a,z,n)}{\partial \ln n} > 0 \quad \forall n$.

Decreasing returns ($\theta < 1$) are standard in an environment of perfect competition and heterogeneous firms. The convex expansion cost reflects information frictions and span-of-control costs of managing a greater number of establishments (Oberfield, Rossi-Hansberg, Sarte, and Trachter, 2024), and the assumption on γ ensures that no firm expands indefinitely.[5] The cost level $\bar{\kappa}$ is the static counterpart of the sunk and continuation costs, $\bar{\kappa}_x$ and $\bar{\kappa}_c$, of the quantitative model. We place no restriction on σ: whether intangible capital and the rival factor are gross substitutes ($\sigma > 1$) or gross complements ($\sigma < 1$) is precisely the question the theory takes to the data.

[5] In Dyrda, Hong, and Steinberg (2026) we develop a theory of a firm, in which the scope and firm's expansion is not limited by a span of control and expansion costs but rather by organizational complexity; a cumulative heterogeneity of firm's footprint.

2.2 The Scope of Production and Intangible Intensity

With the functional form assumptions above, the firm's first-order conditions deliver closed-form expressions for the two endogenous objects of interest, the scope of production and intangible intensity:

$$n = \left(\frac{1}{\gamma \bar{\kappa}} \tilde{\pi}\left(a, z\right) \right)^{\frac{1}{\gamma - 1}} \tag{3}$$

$$s^z = \left[\frac{1-\mu}{\mu} \left(\frac{W}{P} \right)^{1-\sigma} (1+n)^{1-\sigma} + 1 \right]^{-1}. \tag{4}$$

Equation (3) determines scope: the firm expands until the marginal cost of an additional country, which grows at rate $\gamma - 1$ under the convex cost function, equals the per-location profit $\tilde{\pi}(a, z)$. Scope therefore rises with anything that raises per-location profits, including the firm's intangible stock, since the non-rival input raises output in every location at once.

Equation (4) contains the central implication of non-rivalry: scope enters the firm's cost-minimizing input mix. One more unit of the rival factor raises output in a single location, while one more unit of intangible capital raises output in all $1+n$ locations simultaneously, so expansion raises the return to intangible capital relative to the rival factor. How the input mix responds to this change in relative returns is governed by σ.

Proposition 1 (Non-rivalry and the input mix)**.** *Under Assumption 1, intangible intensity is strictly increasing in the scope of production if and only if intangible capital and the rival input are gross substitutes:* $\partial s^z / \partial n > 0$ *if* $\sigma > 1$, $\partial s^z / \partial n < 0$ *if* $\sigma < 1$, *and* s^z *is invariant to* n *in the Cobb-Douglas case* $\sigma = 1$.

Equation (4) gives rise to a structural estimating equation that we can take to the data to estimate σ. However, the geographic scope of a firm is endogenous: according to equation (3), n rises with per-location profits, so between-firm differences in productivity a, the intangible weight μ, and the prices of the rival input and of intangible capital move scope and intensity together and bias the OLS estimates. Even when using panel data to exploit the within-firm variation in n, the firm-level shocks to these characteristics will affect both n and s^z. Identification requires changes in n that are caused by reasons orthogonal to these factors, and the theory points to the natural shifter: the expansion-cost level $\bar{\kappa}$. This cost affects the scope choice through (3) and enters (4) only through n, the exclusion restriction required of a valid instrument. We construct our instrumental

variables in Section 3 based on this theoretical insight.

We also derive a set of comparative-statics results in Appendix A. Under gross substitutability, reductions in the price of intangible capital or in the cost of multinational expansion raise both intangible intensity and scope (Proposition 2), with endogenous expansion amplifying the intensity response (Lemma 2); the two cost reductions are complementary, so their joint effect is super-additive (Proposition 3); and more productive firms choose larger scope and higher intensity and respond more strongly to either cost reduction (Propositions 4 and 5).

For tractability, the theory deliberately abstracts from several forces that shape the joint evolution of intangible investment and multinational expansion in the data. The model is static, intangible investment carries no adjustment costs, the number of foreign countries is a continuous choice with no sunk costs or uncertainty, and operating abroad is costless after paying the expansion cost. All of these frictions will affect the joint process of multinational expansion and intangible accumulation. Section 4 therefore embeds the mechanism in a dynamic model of heterogeneous firms with intangible accumulation subject to adjustment costs and stochastic costs of opening foreign affiliates, ingredients that let it speak to both the cross-section and the dynamics of the firm decisions.

3 Empirical Estimation of the Elasticity of Substitution

This section takes the structural relationship between intangible intensity and the scope of production, equation (4), to firm-level data. We first describe the data and the measurement of the two objects the theory links, the intangible expenditure share s^z and the number of countries of operation n. We then document that multinational firms are more intangible-intensive than domestic firms, that the gap grows with the number of countries in which a firm produces, and that intensity rises after firms first expand abroad. Finally, we estimate the elasticity of substitution σ, instrumenting geographic scope with two sources of exogenous variation in the barriers to foreign direct investment.

3.1 Data

Our primary dataset is the Compustat North America database, which covers publicly listed U.S. and Canadian firms from 1995 to 2023. We use the Annual Fundamentals dataset to extract consolidated, firm-level balance sheet and income statement information, including total revenue, cost of goods sold, research and development (R&D) expenditure, selling, general and administrative (SG&A) expense, as well as tangible and intangible assets. We exclude firm-year observations with

missing industry codes and those with non-positive or missing values of total revenue.

We construct intangible capital as the sum of balance-sheet intangibles and internally generated knowledge capital from capitalized R&D expenditures, following Peters and Taylor (2017). This measure captures the canonical nonrival inputs emphasized by our theory: patents, proprietary software, product designs, and other forms of codified knowledge that can be scaled globally at zero marginal cost. We exclude organizational capital constructed from SG&A expenses, despite its use in related literature (Eisfeldt and Papanikolaou, 2013), because it encompasses substantial location-specific, rival investments (e.g., local employee training, regional supplier relationships). Our primary outcome variable is intangible intensity, defined as the ratio of intangible capital to total assets:

$$s_{it}^{z} \equiv \frac{\text{Intangible Capital}_{it}}{\text{Intangible Capital}_{it} + \text{Tangible Capital}_{it}}. \tag{5}$$

To measure the scope of multinational production, we link Compustat with firm-level subsidiary disclosures to the U.S. Securities and Exchange Commission (SEC), accessed through Wharton Research Data Services (WRDS). Public firms are required to disclose the country of incorporation and ownership structure for subsidiaries in which they hold at least a 10% ownership stake in their Exhibit 21 filings. Our empirical counterpart of n_{it} is the number of distinct foreign countries in which firm i operates subsidiaries in year t.

3.2 Empirical Facts on Multinational Production and Intangible Intensity

Multinational firms are more intangible-intensive than domestic firms, and the gap has widened over time. In our preferred specification, which conditions on year and industry fixed effects and log revenue, the intangible intensity of MNEs exceeds that of domestic firms by 6 percentage points before 2005, roughly a quarter of average intangible intensity at the start of the sample, and the gap nearly doubles to 11 percentage points thereafter (Table 1, column 3). Controlling for revenue widens rather than shrinks the difference—within size categories, MNEs are even more intangible-intensive—while industry composition accounts for about a third of the raw 9-percentage-point gap (columns 1–2).

The gap also scales with the extent of multinational activity. Restricting the sample to MNEs and replacing the indicator with the log number of foreign countries of operation, a 10 percent increase in the number of countries is associated with 0.3 percentage points higher intangible intensity before 2005 and 0.5 percentage points after (Table 1, column 6); for the median MNE, which operates in 8 countries, adding one country implies an increase of about 0.4 to 0.6 percentage

Table 1: Intangible Intensity and Multinational Production

Dependent Var	Intangible Intensity					
	(1)	(2)	(3)	(4)	(5)	(6)
MNE indicator	0.09***	0.12***	0.06***			
	(0.01)	(0.01)	(0.00)			
× post−2005	0.05***	0.05***	0.05***			
	(0.01)	(0.01)	(0.01)			
Log foreign countries				0.04***	0.06***	0.03***
				(0.00)	(0.00)	(0.00)
× post−2005				0.02***	0.02***	0.02***
				(0.00)	(0.00)	(0.00)
Industry FE	N	N	Y	N	N	Y
Control for log revenue	N	Y	Y	N	Y	Y
Observations	27,318	27,318	27,318	11,925	11,925	11,925

Notes: This table reports OLS estimates of the relationship between intangible intensity and multinational activity. The dependent variable in all columns is intangible intensity, defined in equation (5). "MNE indicator" equals one if the firm reports at least one foreign subsidiary in Exhibit 21 of its SEC 10-K filing. "Log foreign countries" is the natural logarithm of one plus the number of distinct foreign countries in which the firm operates subsidiaries. "Post−2005" is an indicator for observations from 2006 onward. Columns (1)–(3) examine the full sample; columns (4)–(6) restrict the sample to MNEs. Balanced panel of the 942 firms listed in every year 1995–2023, so that changes in the MNE gap are not driven by entry and exit (35 percent of post-2005 firm-years come from firms listed after 2005). All specifications include year fixed effects. Columns (2), (3), (5), and (6) include log total revenue as a control for firm size. Columns (3) and (6) additionally include 2-digit NAICS industry fixed effects. Standard errors, reported in parentheses, are clustered at the 2-digit NAICS industry and year level. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

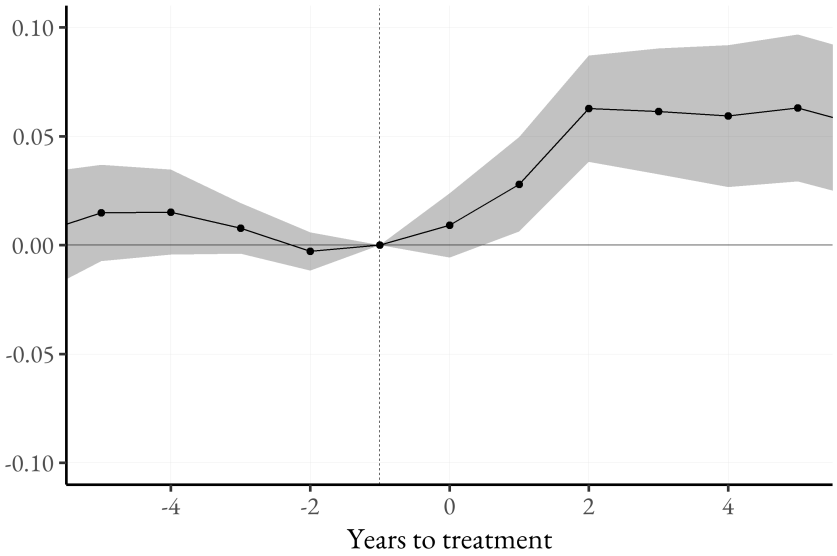
points. Because this gradient holds within multinational firms, it cannot be explained by selection into MNE status alone. It is the raw covariation between s^z and n whose slope, by equation (4), is informative about σ.

Intangible intensity rises within firms after they first become multinational. We estimate the dynamic response of s^z to MNE entry using the interaction-weighted estimator of Sun and Abraham (2021), which is robust to treatment-effect heterogeneity under staggered adoption, on a balanced panel of firms over 1995–2023, with firms that become MNEs later or never as the control group:

$$s^z_{it} = \gamma_i + \delta_t + \sum_{e \notin C} \sum_{\tau \neq -1} \beta_{e,\tau} \cdot \mathbb{1}_{\{E_i = e\}} \cdot \mathbb{1}_{\{t - E_i = \tau\}} + \varepsilon_{it}, \tag{6}$$

where E_i is the year in which firm i first reports a foreign subsidiary. Figure 2 reports the cohort-share-weighted treatment effects. Pre-entry coefficients are close to zero, consistent with parallel trends and no anticipation, and intangible intensity rises by approximately 6 percentage points within 2–4 years of entry, similar in magnitude to the cross-sectional gap. The gradual buildup of the response indicates adjustment costs in intangible investment, an ingredient of the quantitative

Figure 2: Dynamic Effects of MNE Entry on Intangible Intensity



Notes: This figure shows dynamic treatment effects of first attaining MNE status on intangible intensity, estimated from equation (6) with the interaction-weighted estimator of Sun and Abraham (2021). The plotted coefficients are cohort-share-weighted averages of the $\hat{\beta}_{e,\tau}$; the omitted category is $\tau = -1$, one year before entry. The sample is a balanced panel of firms over 1995–2023, and estimates are sales-weighted. Shaded bands represent 95% confidence intervals based on standard errors clustered at the firm and year level.

model in Section 4.

These patterns are what the theory predicts when intangible capital and rival inputs are gross substitutes, but they do not by themselves identify σ: firms choose when and where to expand, so both the timing of entry and the extent of scope may respond to the same profitability shocks that raise intangible investment. The remainder of this section addresses this endogeneity by instrumenting n with policy-driven changes in the barriers to foreign direct investment.

3.3 Estimating the Elasticity of Substitution

Taking logs of the structural equation (4) makes the relationship between intangible intensity and scope linear and delivers our estimating equation:

$$\log\left(\frac{1 - s^z_{it}}{s^z_{it}}\right) = (1 - \sigma)\log(1 + n_{it}) + \alpha_i + \delta_t + \varepsilon_{it}, \tag{7}$$

where i indexes a firm and t indexes year, s^z_{it} is the intangible capital share, and n_{it} is the number of foreign countries in which the firm operates. The two fixed effects absorb the remaining terms of (4). The firm fixed effect α_i absorbs $\sigma \log\left((1 - \mu_i)/\mu_i\right)$, the permanent component of the firm's

intangible weight in the CES technology, along with any other time-invariant firm characteristic. The year fixed effect δ_t absorbs $(1-\sigma)\log(w_t/MC_t)$, the relative price of the rival factor and intangible capital, which is common to all firms within a period, together with aggregate shocks to intangible intensity. The error term ε_{it} therefore absorbs time-varying changes in the intangible weight μ_{it}, and σ is identified from the covariation of intangible intensity and scope that is orthogonal to these changes.

We estimate equation (7) on the unbalanced panel, with σ as the parameter of interest. [6] Since n reflects endogenous multinational expansion decisions that may covary with unobserved determinants of intangible intensity, we instrument for $\log n$ using two sources of exogenous variation in inward FDI barriers across destination countries. Both instruments are constructed as shift-share variables—weighted sums of country-specific policy changes, where the weights are the pre-policy shares of US firms present in each industry.

The first instrument exploits the 2004 EU enlargement, which sharply reduced FDI barriers for firms with pre-existing operations in the eight accession countries. We construct a shift-share instrument whose shares are the leave-one-out fraction of firms in each 3-digit NAICS industry with at least one pre-2004 accession-country subsidiary, and whose shift is the post-2004 indicator, $Z^{\text{EU}}_{ikt} = \frac{n_k^{\text{exp}} - \mathbf{1}[i \text{ exposed}]}{n_k - 1} \times \mathbf{1}[t \geq 2004]$, where n_k is the total number of US firms in industry k and n_k^{exp} is the number of US firms among them that have foreign subsidiaries in the eight accession countries prior to 2004. The identification assumption is that, conditional on firm and year fixed effects, pre-2004 industry-level exposure to accession countries affects intangible intensity only through its effect on geographic expansion n_{it}. A pre-trends test ($p > 0.34$) provides evidence against differential pre-treatment trends between industries with different exposures, supporting the exclusion restriction[7].

The second instrument exploits the staggered expansion of bilateral investment treaty (BIT) networks across host countries. BIT reduces foreign direct investment (FDI) barriers between the two treaty countries by lowering the risks and costs of investing abroad. Letting B_{ct} denote the number of active BITs in country c at time t, we aggregate BIT openness over each industry's pre-period subsidiary locations $Z^{\text{BIT}}_{ikt} = \frac{1}{n_k - 1}[\sum_c n^{\text{pre}}_{kc} \cdot B_{ct} - \sum_c \mathbf{1}[i \text{ in } c \text{ pre-2000}] \cdot B_{ct}]$, where n^{pre}_{kc}

[6] Firm fixed effects absorb permanent differences between entering and exiting firms, and the instrument does not predict exit. The estimate is stable when the sample is restricted to firms observed for at least 5, 10 or 15 years (1.77, 1.79, 1.73) and is 1.97 excluding firm-years near entry or exit.

[7] The pre-trend test interacts each industry's leave-one-out accession-country exposure share with year dummies for 1998–2003 (the pre-treatment period), conditional on firm and year fixed effects. All six pre-period interaction coefficients are individually insignificant (p-values range from 0.34 to 0.64), providing no evidence of differential pre-trends in intangible intensity across industries with varying accession exposure prior to the 2004 enlargement.

is the number of industry-k US firms that have foreign subsidiaries in country c prior to 2000. The BIT instrument relies on the assumption that changes in the number of treaties signed by foreign governments are uncorrelated with the expansion potential of US firms, conditional on firm and year fixed effects. This assumption would be violated if countries sign BITs in response to anticipated FDI demand from the same industries whose firms are in the sample.

Table 2 shows the estimation result of equation (7). The OLS coefficient on $\log n$, Column (1), is -0.038, implying $\hat{\sigma} \approx$1.04, close to Cobb-Douglas. For the IV estimates in Columns (2)–(4), we find that both instruments are positively correlated with $\log n$, although the instrument based on the BIT network exhibits a substantially stronger first stage, with an F-statistic of 18.7 compared to 4.3 for the EU Bartik instrument. The BIT instrument yields an estimate of $\hat{\sigma} = 1.77$. The EU Bartik instrument produces a similar point estimate, $\hat{\sigma} = 1.49$, although this estimate is not statistically distinguishable from 1. Using both instruments jointly results in $\hat{\sigma} = 1.81$. Across specifications, the estimated elasticity of substitution $\hat{\sigma}$ consistently falls between 1.5 and 1.8, well above unity.

Our estimate of $\sigma > 1$ aligns with empirical findings in the recent literature on substitution between intangible and tangible inputs, although the precise definitions of these inputs differ across studies. Aum and Shin (2024) and Jerbashian (2024) estimate elasticities of substitution between software/ICT or IP capital and labor that exceed one.[8] Kleinman (2025) estimates an elasticity of 1.3 between headquarter services—a non-rival input—and branch-level labor within multi-establishment firms, inferred from the higher expenditure shares on headquarter services among firms with broader geographic reach.

Because intangible capital adjusts with a lag, the estimates in Table 2, which use contemporaneous variation within firms, understate the long-run elasticity when σ is greater than one. Incomplete adjustment shrinks the measured response of intensity to scope and pulls $\hat{\sigma}$ toward the Cobb-Douglas value. We therefore read our estimate as a lower bound on the technological elasticity. We will next introduce the quantitative model and illustrate the role of the adjustment cost in the estimation of σ in Section 5. One caveat of the instrument is that the policies could reduce not only the entry cost but also the local operational cost of FDI, the latter of which incentivizes existing establishments to expand in policy areas and may increase their parent firms' intangible intensity without any further geographical expansion. The model will inform the direction and magnitude of such a bias when the policy shifts both the fixed and variable cost of FDI.

[8] Lashkari et al. (2024) estimates the elasticity of substitution between IT and non-IT inputs to be below one. Aum and Shin (2024) notes that the IT capital measured by Lashkari et al. (2024) incorporates both non-rival inputs (e.g., software) with rival inputs (e.g., computers).

Table 2: IV Estimates of the Elasticity of Substitution (σ)

	(1) OLS	(2) EU Bartik	(3) BIT Network	(4) Both
$\log n$	-0.038***	-0.488	-0.773*	-0.810**
	(0.007)	(0.864)	(0.427)	(0.409)
Implied $\hat{\sigma}$	1.04	1.49	1.77	1.81
Firm FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Observations	148,682	148,682	148,682	148,682
Number of firms	16,516	16,516	16,516	16,516
Clustering	Firm	Industry	Industry	Industry
Kleibergen-Paap rk F		4.27	18.67	10.07
Cragg-Donald F		48.99	300.64	143.28
Instruments		Z^{EU}	Z^{BIT}	$\{Z^{\text{EU}}, Z^{\text{BIT}}\}$

Notes: All specifications include firm and year fixed effects. The dependent variable in the second stage is $\log((1 - s_z)/s_z)$ and the endogenous regressor is $\log n$. Standard errors in parentheses. Column (1) reports OLS with standard errors clustered at the firm level. Columns (2)–(4) cluster at the 3-digit NAICS industry level (106 clusters). The EU Bartik instrument uses the leave-one-out share of industry peers with pre-2004 subsidiaries in EU accession countries, interacted with a post-2004 indicator. The BIT network instrument uses the leave-one-out industry-level exposure to host countries' active bilateral investment treaty counts, with subsidiary locations frozen at pre-2000 values. Kleibergen-Paap rk F is the cluster-robust first-stage F-statistic. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4 Quantitative model

This section embeds the mechanism of Section 2—nonrival intangible capital and endogenous multinational scope—into a dynamic, heterogeneous-firm environment. Relative to the static theory, the quantitative model adds life-cycle dynamics, productivity shocks, intangible-capital adjustment costs, an endogenous initial intangible stock, stochastic expansion and contraction, and a per-period cost of continuing foreign operations. These ingredients allow the model to speak to both the cross-sectional and panel dimensions of the data, with multinational status and intangible intensity determined endogenously.

4.1 Environment

There is one home country and N identical foreign countries. The home country is populated by a continuum of firms that produce output using labor, tangible capital, and intangible capital according to a decreasing-returns-to-scale technology. Labor and tangible capital are rival and must be hired separately in each location. As in Section 2, intangible capital is nonrival and can be used simultaneously in all locations. Producing at home is free of fixed multinational costs.

Opening a foreign affiliate requires a sunk cost, while operating an existing affiliate requires a per-period continuation cost. As in Spencer (2022), this distinction follows the tradition of Das, Roberts, and Tybout (2007). Intangible capital is subject to adjustment costs, so firms build it gradually over time, both after entry and after expanding abroad. As in Steinberg (2023) and Ruhl and Willis (2017), aggregate prices are fixed at home and abroad; this is a partial-equilibrium environment.

4.2 Firm characteristics and entry

A firm's state is $(a, z, n) \in \mathcal{S}$, where a is productivity, z is intangible capital, and n is its number of foreign affiliates. Productivity is exogenous and follows a first-order Markov process with transition kernel

$$F(\mathrm{d}a' \mid a) \tag{8}$$

on $\mathcal{A}$. The exact structure of this process, including the fixed and AR(1) components used in the quantitative implementation, is specified in the calibration section. The states z and n are endogenous.

At the end of each period, a firm survives with probability ψ. Firms that do not survive are replaced by entrants with no foreign affiliate. To describe the information available when an entrant chooses its initial intangible stock, let Ω denote the space of entry information and let $\omega \in \Omega$ be distributed according to $H(\mathrm{d}\omega)$. Conditional on ω, the entrant's initial productivity is distributed according to $\bar{F}(\mathrm{d}a \mid \omega)$ on $\mathcal{A}$. Thus ω can reveal persistent components of productivity while leaving a transitory component unresolved. The exact content of ω, and hence the relationship between H, $\bar{F}$, and the incumbent productivity process, is specified in the calibration section. [9]

4.3 Production and static profit maximization

A firm that begins the period with productivity a and intangible capital z produces in each location in which it operates using local labor ℓ and tangible capital k. Define the rival-input bundle as

$$g(k, \ell) = k^{\alpha} \ell^{1-\alpha}. \tag{9}$$

The production technology is

$$f(a, z, k, \ell) = a \left[\mu^{1/\sigma} z^{\frac{\sigma-1}{\sigma}} + (1-\mu)^{1/\sigma} g(k, \ell)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\theta\sigma}{\sigma-1}}. \tag{10}$$

[9] Our Compustat data do not allow us to observe true firm death and creation directly. Firms leave the sample because of delisting and M&A activity, while firms enter when they go public. We therefore do not interpret model entrants as newborn firms that must begin with negligible capital stocks.

The parameter α governs the tangible-capital share in the rival-input bundle, μ controls the importance of intangibles, and σ is the elasticity of substitution between intangible capital and rival inputs. The parameter $\theta < 1$ governs returns to scale at the plant level. Fixed factor-augmenting constants and the common price normalization used in the numerical implementation are suppressed here because they do not change the economic mechanism.

Regulatory and technological barriers to deploying intangible capital abroad are captured by scaling foreign revenue by $\eta < 1$. Given its state (a, z, n), the firm chooses local production inputs to maximize global static profits,

$$\pi(a, z, n) = \max_{\{k_j, \ell_j\}_{j=0}^{n}} \Bigg\{ f(a, z, k_0, \ell_0) - W\ell_0 - Rk_0 + \sum_{j=1}^{n} \left[\eta f(a, z, k_j, \ell_j) - W\ell_j - Rk_j\right] \Bigg\}. \tag{11}$$

Here W is the wage and R is the effective cost of tangible capital, including all rental and depreciation costs, in units of output. Both prices are equalized across home and foreign locations without loss of generality; η captures the effective foreign revenue wedge used in the quantitative implementation. The same intangible stock enters production in every location, whereas labor and tangible capital are chosen separately at each location.

4.4 Multinational expansion and contraction

At the beginning of the period, after observing (a, z, n), the firm first chooses whether to contract, maintain, or expand its affiliate network. This decision has a nested-logit structure. The outer nest is a decision over contraction versus continuation. Conditional on continuation, the inner nest is a decision over remaining at the current affiliate count versus expanding. This structure allows expansion and contraction probabilities to respond to different amounts of idiosyncratic cost variation while retaining closed-form choice probabilities. Opening an additional foreign affiliate requires a sunk cost κ_x; maintaining the current number of affiliates requires a fixed continuation cost κ_c. These costs are iid across firms and over time, drawn from Gumbel distributions with means $\bar{\kappa}_x$ and $\bar{\kappa}_c$ and dispersions ϵ_x and ϵ_c respectively. When $\epsilon_c = \epsilon_x$, the nested choice collapses to a standard three-way logit.

Let $V(a, z, n)$ denote the value after the affiliate decision has resulted in affiliate count n, but

before the firm chooses intangible investment. The value of continuing rather than contracting is

$$V^c(a,z,n) = \epsilon_x \log\left[\exp\left(\frac{V(a,z,n)}{\epsilon_x}\right) + \mathbb{1}_{\{n<N\}}\exp\left(\frac{V(a,z,n+1)-\bar{\kappa}_x}{\epsilon_x}\right)\right]. \tag{12}$$

The value before the contraction decision is

$$\widetilde{V}(a,z,n) = \epsilon_c \log\left[\mathbb{1}_{\{n>0\}}\exp\left(\frac{V(a,z,n-1)}{\epsilon_c}\right) + \exp\left(\frac{V^c(a,z,n)}{\epsilon_c}\right)\right]. \tag{13}$$

An unavailable alternative is assigned value $-\infty$: contraction is unavailable at $n = 0$, and expansion is unavailable at $n = N$.

The probability of contraction is

$$p_-(a,z,n) = \frac{\mathbb{1}_{\{n>0\}}\exp(V(a,z,n-1)/\epsilon_c)}{\mathbb{1}_{\{n>0\}}\exp(V(a,z,n-1)/\epsilon_c) + \exp(V^c(a,z,n)/\epsilon_c)}. \tag{14}$$

Conditional on not contracting, the probability of expansion is

$$p_{+|c}(a,z,n) = \frac{\mathbb{1}_{\{n<N\}}\exp([V(a,z,n+1)-\bar{\kappa}_x]/\epsilon_x)}{\exp(V(a,z,n)/\epsilon_x) + \mathbb{1}_{\{n<N\}}\exp([V(a,z,n+1)-\bar{\kappa}_x]/\epsilon_x)}. \tag{15}$$

The unconditional expansion and stay probabilities are therefore

$$p_+(a,z,n) = [1-p_-(a,z,n)]p_{+|c}(a,z,n), \tag{16}$$

$$p_0(a,z,n) = [1-p_-(a,z,n)][1-p_{+|c}(a,z,n)]. \tag{17}$$

4.5 Intangible capital accumulation and the incumbent's problem

After the affiliate decision, the incumbent chooses gross intangible investment $x \in \mathcal{X}(z)$, denominated in units of the final good. Intangible capital evolves according to

$$z' = (1-\delta)z + x, \tag{18}$$

where $\delta \in (0,1)$ is the depreciation rate. Investment costs

$$Px + \Phi(x,z), \tag{19}$$

where P is the price of intangible capital and Φ is a convex adjustment cost. The parameterization of Φ used in the quantitative implementation is specified in the calibration section.

The post-affiliate-choice value function satisfies

$$V(a,z,n) = \max_{x \in \mathcal{X}(z)} \left\{ \pi(a,z,n) - \kappa_c n - Px - \Phi(x,z) + \beta\psi \int_{\mathcal{A}} \widetilde{V}(a',z',n) F(\mathrm{d}a' \mid a) \right\}, \tag{20}$$

subject to equation (18). The current affiliate choice is therefore made before intangible investment: a firm that expands or contracts chooses x using its new affiliate count. Productivity then evolves, and the next period begins with another affiliate decision. The sunk expansion cost appears inside the expansion alternative in equation (12), and firm death has no continuation or salvage value, so ψ multiplies the continuation value.

4.6 Entrant's problem

An entrant begins with no foreign affiliate and chooses an initial stock before the component of productivity not contained in ω is realized. Its policy function is

$$z_0(\omega) \in \arg\max_{z \in \mathcal{Z}} \left\{ \int_{\mathcal{A}} V(a,z,0)\, \bar{F}(\mathrm{d}a \mid \omega) - Pz \right\}. \tag{21}$$

Thus the entrant chooses the setup stock that maximizes its expected value as a domestic incumbent, conditional on the information available at entry, net of the cost of purchasing that stock. The initial purchase is a setup decision and does not incur the incumbent adjustment cost Φ. In the quantitative implementation, ω contains the entrant's fixed characteristics but not the initial realization of the transitory productivity component; the precise components are specified in the calibration section.

4.7 Aggregation and stationary equilibrium

Let $\mathcal{A} \subset \mathbb{R}_{++}$ and $\mathcal{Z} \subset \mathbb{R}_{++}$ denote the support of productivity and intangible capital, respectively, and let and $\mathcal{N} = \{0, \ldots, N\}$ denote the set of feasible affiliate counts. From affiliate count n, the feasible changes are

$$\mathcal{D}(n) = \{d \in \{-1, 0, +1\} : n + d \in \mathcal{N}\}. \tag{22}$$

For $d \in \mathcal{D}(n)$, let $p_d(a,z,n)$ denote the corresponding contraction, stay, or expansion probability.

Let Ψ denote the distribution of end-of-period states of the state space $\mathcal{S} = \mathcal{A} \times \mathcal{Z} \times \mathcal{N}$. For

any measurable $B \subseteq \mathcal{S}$, the transition kernel is

$$Q((a,z,n),B) = \psi \int_{\mathcal{A}} \sum_{d \in \mathcal{D}(n)} p_d(a',z,n) \mathbb{1}_{\{(a',g_d(a',z,n),n+d) \in B\}} F(\mathrm{d}a' \mid a) + (1-\psi) \int_{\Omega} \int_{\mathcal{A}} \mathbb{1}_{\{(a_0,z_0(\omega),0) \in B\}} \bar{F}(\mathrm{d}a_0 \mid \omega) H(\mathrm{d}\omega), \tag{23}$$

where $g_d(a,z,n) = (1-\delta)z + x(a,z,n+d)$. The first line describes surviving firms: productivity transitions before the affiliate and investment policies are evaluated, and investment uses the post-choice affiliate count. The second line describes death and replacement by entrants. The stationary distribution satisfies

$$\Psi(B) = \int_{\mathcal{S}} Q(s,B)\, \Psi(\mathrm{d}s) \qquad \text{for every measurable } B \subseteq \mathcal{S}. \tag{24}$$

Given exogenous prices, a stationary equilibrium consists of value functions V and $\widetilde{V}$, an incumbent intangible-investment policy x, affiliate-choice probabilities (p_-, p_0, p_+), an entrant policy z_0, and a stationary distribution Ψ satisfying equations (20), (21), and (24).

5 Structural estimation

We estimate the parameters of the quantitative model to match cross-sectional and panel moments from the Compustat sample described in section 3.1 above. Our approach borrows from Duarte and Fonseca (2026). We solve the model once for many parameter vectors simultaneously in parallel, train a neural network on the resulting simulated datasets to approximate the mapping between parameters and moments, and estimate the parameter vector that minimizes the error between the approximated model and the data.

This approach has three advantages over traditional simulated-method-of-moments strategies. First, we only need to run a small number of fixed "batches" of model solves to construct the survey, whereas traditional methods require new solves every time the candidate solution vector (for local optimization algorithms) or the population (for genetic algorithms) is updated. This allows us to make better use of parallel computing; an estimation using our approach that requires the same total number of model solves as a traditional structural estimation will require far less wall time. Second, the neural network is smooth, which means that it—and its Jacobian, which is computed using automatic differentiation—can be evaluated at essentially zero cost. This means that once the network has been trained, finding its minimum using many starting values is almost instantaneous,

and identification can easily be assessed both globally and locally. Third, the separation between the actual model solves and evaluations of the network surrogate allows us to experiment with different weighting matrices, numerical-optimization algorithms, and hyperparameters of those algorithms (tolerances, step sizes, etc.) without having to either resolve the model or retrain the network.

5.1 Parameterizations, estimated parameters, and moment selection

The firm's productivity, a, is specified as the product of two components: a fixed effect a_1 that is drawn from a truncated Pareto distribution with tail parameter γ_{a_1} and is constant across a firm's life; and a_2, which follows an AR(1) process in logs, $\log a_{2,t+1} = \rho_{a_2} \log a_{2,t} + \varepsilon_{t+1}$, where ε_{t+1} is normally distributed with mean zero and variance $\sigma^2_{a_2}$. We assume that entrants observe the fixed effect a_1 when choosing their initial intangible capital but not the initial stochastic component a_2, which is subsequently drawn from its stationary distribution. The intangible-capital adjustment cost takes the standard quadratic form, $\Phi(x, z) = \frac{\phi}{2} z \left(\frac{x}{z} - \delta_z\right)^2$, where δ_z is the depreciation rate of intangible capital.

Several parameters are assigned externally. The number of foreign countries N is set to one; firms can simply either be non-MNEs or MNEs. The wage, W, and the cost of intangible capital, P, are set to one without loss of generality. The cost of tangible capital, R, is set to 12 percent, representing an interest rate of 2 percent and a depreciation rate of 10 percent. The share of tangible capital in the rival bundle, α, is set to 5 percent; this parameter governs the relative use of tangible capital and labor within the rival bundle, and its value is inherited from our existing calibration rather than separately disciplined by the moments below. The depreciation rate of intangible capital, δ_z, is set to 27.5 percent, in the range of estimated depreciation rates for R&D capital (Li and Hall, 2016). The firm's discount factor, β, is set to 0.95. The degree of decreasing returns to scale, θ, is set to 0.7.

Two parameters are set directly based on moments in the Compustat data, as they have one-for-one counterparts. The survival rate, ψ, is set equal to one minus the sample attrition rate of 7.7 percent. The fixed effect's Pareto tail, γ_{a_1}, is set using the empirical sales-tail slope. Optimal sales scale with fixed productivity as $a_1^{1/(1-\theta)}$, so a Pareto tail exponent γ_{a_1} for fixed productivity implies a sales-tail exponent $\gamma_{a_1}(1-\theta)$. The slope in the data is 1.0633, so we set $\gamma_{a_1} = 1.0633/(1-\theta) = 3.5443$.

We fix the mean continuation cost, $\bar{\kappa}_c$, at zero, after several estimations pinned it at this lower bound; a negative mean cost did not make economic sense to us. The continuation-cost dispersion,

ϵ_{κ_c}, remains free. This leaves nine parameters that must be jointly estimated, whose symbols and descriptions are listed in the first two columns of Table 3 below. We estimate these parameters to match fifteen moments from the Compustat data, whose descriptions, values, and bootstrapped standard errors are shown in the first three columns of Table 4. The first five are cross-sectional moments that capture variation across firms. The remaining ten are panel moments that capture variation within firms over time. Event-study moments are not targeted in the reported estimation.

Table 3: Estimated parameters, survey bounds, and identification measures

Parameter	Description	Survey bounds	Estimate	Sandwich SE	Fold SD
σ	Intangible–rival elasticity	$[1.500,\ 3.500]$	2.250	0.3237	0.07801
μ	Intangible share in production	$[0.010,\ 0.080]$	0.05010	0.01493	0.002901
ρ_{a_2}	Stochastic productivity persistence	$[0.750,\ 0.925]$	0.8855	0.01315	0.001055
$\sigma^2_{a_2}$	Stochastic productivity variance	$[0.0025,\ 0.090]$	0.02160	0.001474	0.0001276
ϕ	Adjustment-cost scale	$[0.010,\ 5.000]$	0.2132	0.1451	0.02179
η	Variable FDI barrier	$[0.400,\ 0.900]$	0.6628	0.01330	0.001814
$\bar{\kappa}_x$	Mean entry cost	$[0.050,\ 0.500]$	0.3374	0.03868	0.01020
ϵ_{κ_x}	Entry-cost dispersion	$[0.020,\ 0.160]$	0.1104	0.01390	0.003526
ϵ_{κ_c}	Continuation-cost dispersion	$[0.0010,\ 0.020]$	0.008380	0.001109	0.0002106

Notes: ϕ is sampled in $\log_{10}$ units but reported in levels; its sandwich standard error uses the delta method, and its fold SD is computed after transforming each fold estimate to levels. $\bar{\kappa}_c$ is fixed at zero. $\sigma^2_{a_2}$ is estimated as a standard deviation but reported as a variance, with uncertainty transformed accordingly. The sandwich SE measures sampling error in the empirical moments; the fold SD measures surrogate-training variation.

5.2 Stage 1: Surveying the parameter space

The first stage of our estimation procedure is a survey of the model's behavior across the full range of plausible parameter values. We construct a nine-dimensional box in the parameter space using the bounds shown in column three of Table 3. We then draw a Sobol sequence of $N_1 = 32,768$ parameter vectors from this box to ensure uniform coverage in all regions. The adjustment-cost parameter, ϕ, is drawn using log-10 spacing; the other eight coordinates are spaced linearly, including the standard deviation σ_{a_2} rather than its reported variance.

For each of these vectors, we solve the model and simulate a panel dataset with 50,000 firms and 100 recorded periods after a 50-period burn-in. We then compute the model counterparts of the target moments.[10] This step is embarrassingly parallel—each parameter vector can be evalu-

[10]For points that are degenerate, in the sense that some of the moments cannot be computed (because, e.g., multinational participation is zero), we set these moments to missing values rather than zeros. This ensures that the next step trains the neural network only on points where the moments are well-defined. Solver convergence is checked separately from moment availability. The model is stationary and has a single industry, so its moments are computed directly from the simulated panel rather than by estimating the empirical industry and year fixed effects.

Table 4: Targeted moments and model fit

Moment	Data	Bootstrap SE	Weight	Network approximation	Ground truth
Cross-sectional moments					
Share of firms that are MNEs	0.3528	0.0044	1	0.3533	0.3525
Average intangible intensity	0.4255	0.0034	1	0.4254	0.4255
Intangible-intensity gap, MNEs minus non-MNEs	0.0513	0.0045	1	0.0497	0.0492
Foreign-to-domestic sales ratio for MNEs	0.2596	0.0133	1	0.2583	0.2565
Slope of log sales on intangible intensity, age 5+	0.3984	0.1176	5	0.4214	0.4316
Panel moments					
SD of one-period changes in log sales	0.4695	0.0062	0.5	0.4706	0.4700
Average within-firm SD of log sales	0.5665	0.0058	0.5	0.5673	0.5662
Variance of three-period changes in log sales	0.6078	0.0158	0.34	0.5803	0.5789
Variance of five-period changes in log sales	0.8511	0.0232	0.34	0.8856	0.8807
Variance of seven-period changes in log sales	1.0323	0.0333	0.34	1.1422	1.1363
Autocorrelation of intangible intensity	0.9302	0.0018	0.5	0.8312	0.8312
Median SD of one-period investment-rate changes	0.1140	0.0019	1	0.1130	0.1109
MNE entry rate, two-year confirmed	0.0529	0.0011	1	0.0528	0.0529
MNE exit rate, two-year confirmed	0.0565	0.0016	1	0.0564	0.0560
MNE exit hazard after five or more periods	0.0328	0.0012	0.5	0.0328	0.0329

Notes: Bootstrap standard errors resample firms in blocks. Weight is the economic tier weight w_n in the minimum-distance objective. Network approximation is the output of the trained neural network for a given moment evaluated at the estimated parameter vector $\hat{\theta}$. Ground truth is the moment obtained by re-solving and simulating the full structural model at $\hat{\theta}$.

ated independently, and there is no need for communication between compute nodes.[11] The output is a dataset with N_1 rows. The portion used in estimation has 24 parameter-and-target-moment columns (nine parameters plus fifteen moments); the files also retain additional untargeted moments, metadata, and computational diagnostics. This dataset depicts the mapping from parameters to moments over the global box. It can be reused when changing target values, weights, or network architecture, provided the structural model and simulated-moment definitions remain unchanged and the search stays within the surveyed bounds.

5.3 Neural-network approximation

We train a neural network (or really, a set of separate networks, one for each simulated moment, including each of the 15 target moments) to learn to approximate this mapping. For each target moment m_n for $n = 1, \ldots, 15$, the task is to predict that moment's value from the parameter vector, using the survey from stage 1 as fitting data.

[11] We run the survey on the University of Toronto's Trillium cluster.

The architecture we choose is a multilayer perceptron (MLP), which is a composition of alternating linear maps and a nonlinear activation function. The MLP can be represented mathematically as

$$h_0 = u, \quad h_\ell = \tanh\left(W_\ell h_{\ell-1} + b_\ell\right) \text{ for } \ell = 1, \ldots, L, \quad \hat{m}_n(\theta) = w_{\text{out}}^\top h_L + b_{\text{out}}, \tag{25}$$

where θ is a parameter vector, u is that vector after it has been rescaled to the unit hypercube, and $\hat{m}_n$ is the predicted value of target moment n. The MLP passes u through L hidden layers, each of which first applies an affine map $h \mapsto W_\ell h + b_\ell$ and then applies a nonlinear activation function (which in our case is the hyperbolic tangent), before passing the result through a final affine output layer $w_{\text{out}}^\top h + b_{\text{out}}$. We set the number of layers L to 3 and the dimension of each layer to 64. We use $\tanh$ as the activation function, rather than kinked alternatives such as the rectified linear unit (ReLU), which are common in machine learning, because it guarantees continuous differentiability, which is essential to the next stage of the procedure. The matrices, W_ℓ, and vectors, $b_\ell, w_{\text{out}}, b_{\text{out}}$, are the network parameters that we need to solve for.

We choose these weights to minimize the mean squared prediction error over the survey dataset, first using the global survey and then the pooled survey described below. We use a standard stochastic gradient descent (SGD) approach: starting from random weights, we repeatedly move in the direction that most reduces the error, where each step's direction uses the gradient from a random subset (a "mini-batch") of rows.[12] To guard against overfitting, we use cross-validation: we partition the training survey into K approximately equally-sized partitions ("folds"), and train $K = 10$ different versions of the moment networks, each of which excludes the survey data from one of the folds. Then, to actually make predictions and estimate parameters, we use the ensemble network $\hat{m}_n(\theta) = \frac{1}{K}\sum_{k=1}^{K} \hat{m}_n^k(\theta)$, where $\hat{m}_n^k(\theta)$ represents the fold-k network (the one that held out fold k). This ensemble network is essentially free to evaluate—predicting all fifteen moments from a given parameter vector is just a handful of small matrix operations—and its partial derivatives can be computed exactly using automatic differentiation across the entire survey space.

The cross-validation approach serves two purposes. First, since each fold-k network has some data it never saw during training, that data can be used to measure out-of-sample fit. Specifically, we compute the out-of-fold R^2, which measures the share of a moment's variance across the survey that the network predicts on points outside its training set. In the final pooled survey, the lowest mean out-of-fold R^2 across the target moments is 0.926, for investment-rate volatility; it is

[12] We use the *Adam* optimizer, which is an adaptive SGD variant that rescales step sizes using a running average of past gradients as well as the current gradient, for 40,000 training steps.

0.977 for the sales–intensity slope. These statistics summarize prediction over the survey distribution, not necessarily accuracy in the neighborhood selected by the optimizer. Second, since the different fold-k networks constitute slightly different versions of the mapping between parameters and moments, the spread between their predictions provides a measure of uncertainty about the mapping's accuracy, which we use to provide a measure of parameter estimation error below.

5.4 On-network estimation and refinement

With the network in hand, we estimate the parameter vector that minimizes the weighted sum of the squared errors between the approximated model and the data,

$$L(\theta) = [\hat{m}(\theta) - m]' \, W \, [\hat{m}(\theta) - m], \quad W = \operatorname{diag}\left(\frac{w_n}{s_n^2}\right), \tag{26}$$

where w_n is the economic weight of moment n (a hyperparameter we choose ourselves), and s_n is its residual scale, which we set to the maximum of the moment's empirical bootstrap standard error and two percent of its value: $s_n = \max\{\mathrm{SE}^{\mathrm{ref}}(m_n), 0.02|m_n^{\mathrm{ref}}|\}$. We choose not to weight by the inverse of the empirical covariance matrix (which maximizes efficiency in standard GMM) for three reasons. First, this weighting scheme can be badly behaved in samples of this size (Altonji and Segal, 1996). Second, some of the moment blocks are near-singular, so the inverse is numerically delicate. Third, and most important, a precisely-measured moment that the model cannot structurally reproduce would dominate the fit and force every parameter to help accommodate that feature of the data; the economic weights w_n help us sidestep this problem. The chosen weights are shown in column four of Table 4.

We minimize $L(\theta)$ using a multistart optimization approach. We draw 512 separate scrambled Sobol points in the survey space, use a gradient-based local optimization algorithm to minimize the objective starting from each point, and select the best estimate across all of these runs. Since the network can be evaluated instantaneously and these runs can be parallelized, this stage of the procedure requires very few computational resources.

After obtaining a preliminary estimate of the parameter vector, we repeat the survey, network approximation, and estimation stages one more time. Our second survey refines coverage around all of the candidate optima from the ten separate k-fold networks in addition to the ensemble network's optimum. We draw a second set of 32,768 Sobol points, allocated approximately equally across eleven boxes centered at these estimates. Each box has a half-width equal to 10 percent of the corresponding original range. After solving and simulating these new points, we pool them with

the original survey to obtain 65,536 parameter vectors, retrain the moment networks on the larger dataset, and re-estimate the minimizing parameter vector. This design increases density around the range of optima suggested by the initial networks without abandoning global coverage.

5.5 Estimation results

The fourth column of Table 3 shows our estimated parameter values after the refinement step. The headline result is our structural estimate of the elasticity of substitution between intangible capital and rival factors, σ, which is 2.25, about 25 percent larger than our reduced-form estimate. The persistence of stochastic productivity, ρ_{a_2}, is 0.89 and its innovation variance, $\sigma^2_{a_2}$, is 0.022. The estimated foreign-revenue multiplier, η, is 0.66. The mean cost of setting up foreign affiliates, $\bar{\kappa}_x$, is 0.34, while the mean continuation cost is fixed at zero. Entry-cost dispersion is 0.11, about thirteen times the continuation-cost dispersion of 0.008. The positive entry cost is consistent with sunk costs of foreign participation.[13]

The last two columns of Table 4 report the fit of the estimated model against the data. The fifth column (network approximation) shows the prediction of the neural-network approximation of the model at the estimated parameter vector, and the last column shows the ground truth, i.e., the result obtained by solving the full quantitative model at that vector. The ground truth is generally close to the approximation at this evaluated point. For the sales–intensity slope, which plays a crucial role in identification as we document below, the network predicts 0.42 and the full model delivers 0.43, compared with 0.40 in the data. Most empirical targets are closely matched. The principal shortfall is intangible-intensity persistence, whose autocorrelation is 0.83 in the model versus 0.93 in the data. The model also overpredicts the variance of seven-period sales changes, 1.14 versus 1.03.

5.6 Identification

We use three sets of tools to assess the identification of our estimated parameters.

5.6.1 Global identification

The neural network approximation of the mapping between parameters and moments can be used to characterize how sensitive the objective function is to each parameter along the entirety of its bounding box. The diagnostic tool we use is a simple visualization: the minimum-loss profile.

[13] See, e.g., Das et al. (2007) and Alessandria and Choi (2007) for exporting and Spencer (2021) for multinational production.

For each parameter j, we hold that parameter fixed at a particular value v in the survey space (not necessarily the actual optimum), re-optimize all of the other parameters, and record the best attainable value of the loss function:

$$L_j(v) = \min_{\theta:\theta_j=v} L(\theta). \tag{27}$$

By varying v, we can construct a curve that illustrates how sensitive the fit is to changes in a given parameter. A deep, sharp valley indicates good identification: the fit deteriorates when a parameter moves away from its estimate, even when all the other parameters can move to compensate. A flat profile indicates poor identification because the target moments are indifferent to that parameter.

One important note: the estimated value of a parameter need not line up exactly with the lowest point drawn on its minimum-loss profile. The estimate is obtained by optimizing all parameters continuously, whereas the profile is plotted only at a finite grid of fixed values for the parameter in question and separately re-optimizes the remaining parameters at each grid point. The estimate may therefore lie between two plotted grid points, especially when the bottom is relatively flat. In the continuous limit, however, the profile evaluated at the estimated value must attain the same loss as the joint estimate, provided both optimizations find the same basin.

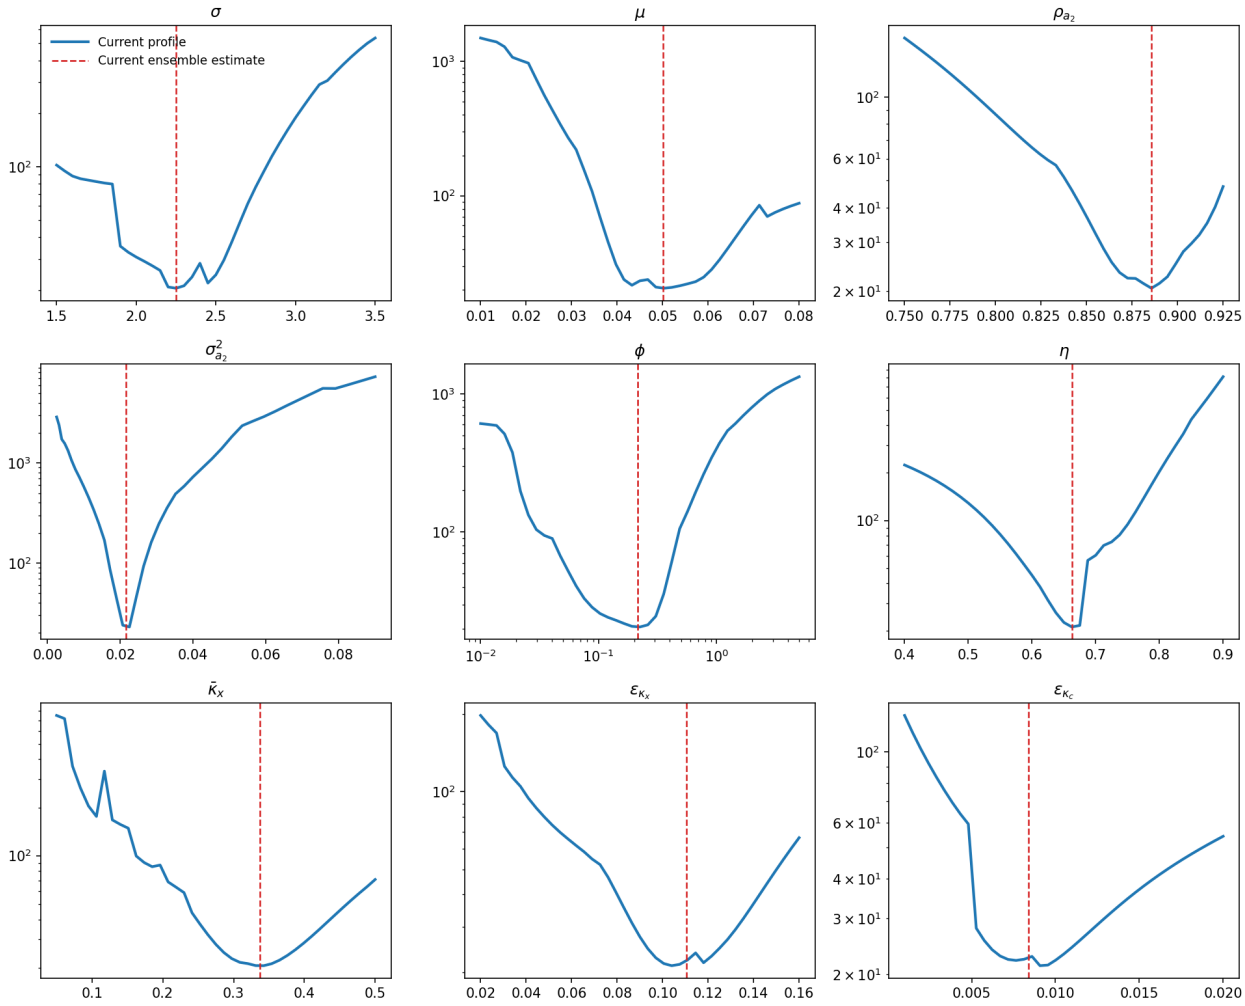
Figure 3 shows the minimum-loss profiles. All nine estimates are interior to their survey bounds, although the depth and shape of their basins vary. The key factor substitution elasticity, σ, lies in a clear valley, while the profile for stochastic productivity variance, $\sigma^2_{a_2}$, rises particularly steeply away from its minimum. The profiles for ρ_{a_2} and entry-cost dispersion, ϵ_{κ_x}, are comparatively shallow toward their upper bounds. Continuation-cost dispersion, ϵ_{κ_c}, also has a broad upper-side profile. The valley for mean entry cost, $\bar{\kappa}_x$, is globally deep but locally relatively flat. These are diagnostics of the surrogate objective within the surveyed box, not confidence intervals or a proof of global identification. In particular, they are conditional on fixing $\bar{\kappa}_c = 0$; there is no profile for that fixed parameter.

5.6.2 Local identification

We also report two measures of how well the parameters are locally identified near the optimum, each of which answers a different question. The first is the sandwich standard error, which measures how much $\hat{\theta}$ would move across hypothetical resamples of the Compustat panel. Its formula is

$$\widehat{\mathrm{Var}}(\hat{\theta}) = (G'WG)^{-1}G'WSWG(G'WG)^{-1}, \tag{28}$$

Figure 3: Minimum-loss profiles for all estimated parameters



Notes: Each panel shows the minimum-loss profile and estimated value for one of the internally estimated parameters. Blue lines: minimum-loss profiles. Red dashed vertical lines: estimated parameter values. X-axis: range of surveyed values for the parameter in question. Y-axis: best attainable value of the loss function holding fixed that parameter at a value v, $L_j(v) = \min_{\theta:\theta_j=v} L(\theta)$.

where G is the Jacobian matrix of the partial derivatives of the approximated mapping between parameters and moments, S is the joint firm-bootstrap covariance matrix of the data moments, and W is the weighting matrix used in estimation, including its fixed reference residual scales.[14] The results are shown in column 5 of Table 3. The sandwich standard error for the key elasticity, σ, is 0.3237—about 14.4 percent of its estimated value. The mean entry cost and its dispersion are about 8.7 and 7.9 times their standard errors, respectively. Precision is less uniform than these ratios suggest: the standard error for ϕ is 0.1451, about 68 percent of its estimate, and that for μ is 0.01493, about 30 percent. Thus the local linearization indicates meaningful sampling uncertainty for some parameters, rather than uniformly tight identification.

The second measure is the cross-fold standard deviation, which is a measure of disagreement between the different k-fold neural networks—how much the estimate moves when the map is redrawn using a different nine-tenths of the training data. It is constructed by re-running the estimation separately on each of these networks instead of using the ensemble, and then measuring the standard deviation of the resulting parameter estimates. The cross-fold standard deviations are smaller than the sandwich standard errors for all nine parameters. For σ, the cross-fold standard deviation is 0.07801, about 3.5 percent of its estimated value, compared with a sandwich standard error of 0.3237. This indicates greater stability across the fitted networks than the empirical-moment sampling uncertainty would imply, although it cannot rule out approximation errors shared by all folds. There is no conventional t-statistic or critical value for this measure: the folds overlap and the fold estimates are not independent draws from a sampling distribution. We therefore treat it as a surrogate-stability diagnostic and interpret it relative to the parameter estimate, the sandwich standard error, and the width of the minimum-loss profile.

5.6.3 Sensitivity to target moments

To analyze the influence of each target moment in identifying the parameters, we use the Andrews, Gentzkow, and Shapiro (2017) sensitivity diagnostic, which is defined as

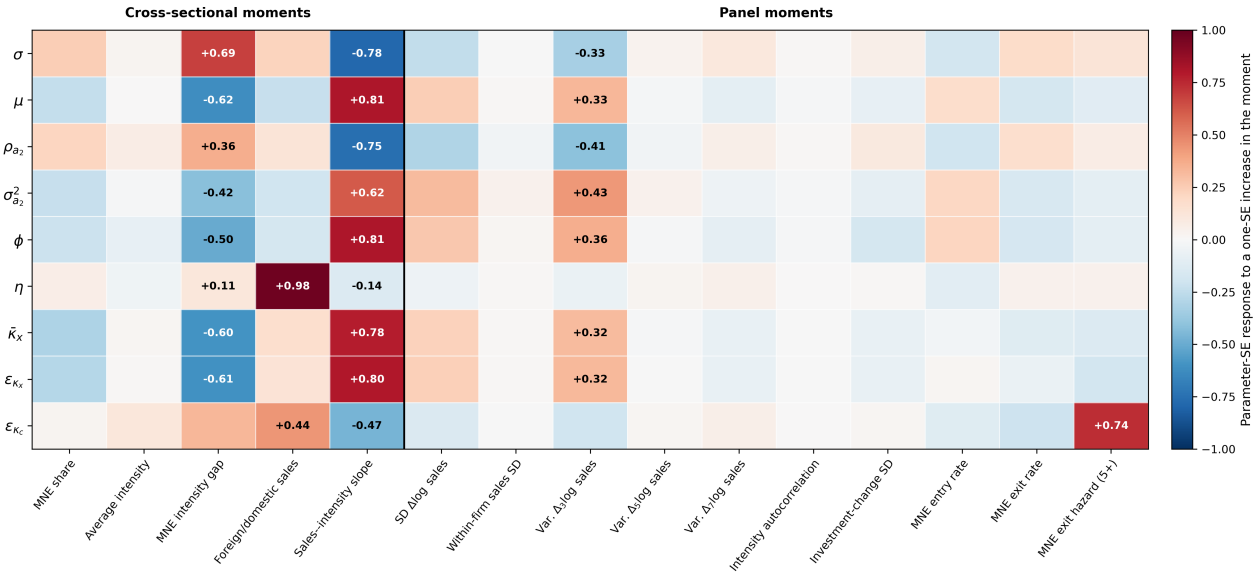
$$\Lambda_{\mathrm{AGS}} = -\left(G'WG\right)^{-1} G'W, \qquad D_m \equiv -\Lambda_{\mathrm{AGS}} = \left(G'WG\right)^{-1} G'W \approx \frac{\partial \hat{\theta}}{\partial m}, \tag{29}$$

[14] For consistency, we use this weighting matrix rather than the statistically-efficient inverse covariance matrix. Note that G is often expensive to compute in traditional SMM, as it requires one to perturb each parameter and resolve the model. Automatic differentiation gives exact derivatives of the neural-network surrogate, not of the structural model. We have not repeated a full-model finite-difference Jacobian audit at this estimate. We therefore interpret the sandwich formula as a local linearization diagnostic; it does not incorporate surrogate error, simulation error, uncertainty in externally assigned parameters, or uncertainty about the restriction $\bar{\kappa}_c = 0$. Because the moments are not exactly matched, it also omits residual-weighted second-derivative terms in the objective's curvature.

where, again, G is the Jacobian matrix of the mapping between parameters and moments and W is the user-specified weighting matrix. The derivative interpretation uses the same local linearization as the sandwich calculation; it is not an exact comparative static of the nonlinear objective when residuals are nonzero. The Andrews–Gentzkow–Shapiro measure Λ_{AGS} describes sensitivity to an additive perturbation of the moment condition $\hat{m}(\theta) - m$. Because our question instead raises the empirical target m, we report its negative, D_m: "if target moment n were perturbed upward, how much (and in which direction) would the estimate for parameter j move?"

Figure 4 shows the results of this calculation. Red cells indicate positive relationships between moments and parameters; blue cells indicate negative relationships; deeper colors indicate stronger relationships. To facilitate interpretation, we report $(D_m)_{jn}\,\mathrm{SE}(m_n)/\,\mathrm{SE}(\hat{\theta}_j)$, scaling the perturbation of each moment by its bootstrap standard error and the resulting change in each parameter estimate by its sandwich standard error. This standardization is our reporting normalization, not part of the definition of the raw Andrews–Gentzkow–Shapiro sensitivity measure. It puts the perturbations in units of empirical moment uncertainty and the parameter responses in units of estimated parameter uncertainty.

Figure 4: Andrews et al. (2017) sensitivity of parameter estimates to targeted moments



Notes: Each cell shows the effect of increasing the target moment indicated on the x-axis on the estimated value of the parameter indicated on the y-axis. The size of the increase in the moment is one bootstrap standard deviation. The effect on the parameter is measured in sandwich standard errors. For example, a one-SD increase in the sales–intensity slope decreases the estimated value of σ by 0.78 standard errors.

The key elasticity σ is most sensitive to the sales–intensity slope and the intangible-intensity

gap between MNEs and non-MNEs. Increasing the slope target by one bootstrap standard error lowers the estimated elasticity by 0.78 sandwich standard errors, while increasing the gap target raises it by 0.69 standard errors. The next-largest sensitivity is to the variance of three-period sales changes, with a standardized effect of -0.33. These responses describe joint reoptimization of all nine free parameters. In particular, the negative response to the slope target is not the direct partial effect of σ on the model's sales–intensity slope, and it need not have the same sign.

Several moments play important roles in identifying multiple parameters. The sales–intensity slope has the largest standardized sensitivity for seven of the nine parameters: σ, μ, ρ_{a_2}, $\sigma^2_{a_2}$, ϕ, $\bar{\kappa}_x$, and ϵ_{κ_x}. The MNE intensity gap and three-period sales-change variance also feature prominently in those rows. By contrast, the foreign/domestic sales ratio dominates the sensitivity of the variable FDI barrier, η, with a standardized effect of 0.98. The exit hazard after five or more periods has the largest effect on continuation-cost dispersion, ϵ_{κ_c}, at 0.74, followed by the sales–intensity slope and foreign/domestic sales ratio. No sensitivity is reported for mean continuation cost because it is fixed. These patterns depend on the selected targets, their weights, and the empirical covariance used for standardization; they are not a weight-free ranking of the information in each moment. The overlap across rows makes clear that the free parameters are jointly identified rather than each being pinned down by a single moment.

5.7 IV estimation on simulated data

We use the estimated model to evaluate the performance of the reduced-form instrumental-variables strategy used in section 3.3: whether or not this strategy correctly recovers the structural object, how the answer depends on the specifics of the shock to FDI barriers, and how it depends on the horizon over which firms adjust.

5.7.1 Experiment design

We start with a sample of incumbent firms from the stationary distribution and divide them into four equal-sized groups. Three of these groups receive small, medium, and large shocks; the fourth group is the control and receives no shock. In the first set of experiments, the shock is to the fixed recurring cost, $\bar{\kappa}_c$. The second set of experiments shocks the variable foreign production wedge, η. The third set combines both types of shocks.

Each experiment is an unanticipated, permanent event. Firms in each group observe their assigned policy shock and re-optimize their policy functions for intangible investment and expansion/contraction. Productivity continues to evolve according to the estimated process, but the new

policy functions are used to create the post-shock observations in the simulated dataset. As in the data, we track a balanced panel of incumbent firms conditional on surviving for a certain number of periods. Using this panel, we estimate the analogue of the empirical specification,

$$\Delta \log\left(\frac{1-s_{it}^{z}}{s_{it}^{z}}\right) = (1-\sigma)\,\Delta \log(1+n_{it}) + u_{it}, \tag{30}$$

where we instrument the change in multinational scope with a firm's assigned group (control, treated-low, treated-medium, or treated-high). We report short- and long run estimates that use data from one and twenty periods after the shock, respectively, to investigate the role of adjustment frictions.

Rather than using four groups of different firms, we draw one set of 200,000 incumbent firms from the stationary distribution and initialize the four groups as copies of this set. The four copies of a given firm start from identical states and receive identical idiosyncratic shocks, but follow the policy functions for their own assigned treatment. This matched design reduces Monte Carlo variation from differences in initial states and subsequent idiosyncratic shocks. The copies are not independent observations, so we cluster standard errors cluster by underlying baseline firm. The table reports means and standard deviations across five different versions of each experiment that use different randomization seeds (and thus different idiosyncratic shocks for each firm and its copies). These standard deviations measure simulation variability rather than sampling uncertainty in the empirical estimates; we could artificially reduce conventional standard errors to zero simply by simulating more firms.

5.7.2 Fixed cost reductions

Panel (a) of Table 5 reports the results of our first set of experiments, where firms experience shocks to the fixed recurring multinational production cost, $\bar{\kappa}_c$. Firms in the three treated groups experience an increase in this cost of $m_\kappa + \{0.00025, 0.00050, 0.00075\}$ for $m_\kappa \in \{0.001, 0.0015, 0.002\}$. Each value of m_k is a different experiment, allowing us to explore whether the scale of the shock affects the outcome of the estimation.

The short-run IV estimate is $\widehat{\sigma}_1 = 1.95$, which is about 13 percent below the structural estimate and quite similar to our reduced-form estimates from section 3 (well within these estimates' confidence intervals), whereas the long-run estimate $\widehat{\sigma}_{20}$ is very close to the true value of 2.25. The difference is due to the fact that intangible capital is a state variable that must be accumulated gradually over time due to adjustment costs. The short-run estimate $\widehat{\sigma}_1$ does not fully reflect reflects

Table 5: IV estimates of σ in simulated model data

Shock offset or mixture weight	$\widehat{\sigma}_1$	$\widehat{\sigma}_{20}$
(a) Recurring-cost increase, $\bar{\kappa}_c$		
$m_\kappa = 0.0010$	1.964 (0.029)	2.270 (0.013)
$m_\kappa = 0.0015$	1.957 (0.023)	2.288 (0.012)
$m_\kappa = 0.0020$	1.945 (0.022)	2.306 (0.012)
(b) Foreign-revenue multiplier reduction, η		
$m_\eta = 0.050$	1.600 (0.020)	2.872 (0.013)
$m_\eta = 0.075$	1.628 (0.017)	2.910 (0.012)
$m_\eta = 0.100$	1.661 (0.016)	2.984 (0.013)
(c) Combination of the middle cases		
$\lambda = 0.25$	1.816 (0.020)	2.559 (0.012)
$\lambda = 0.50$	1.715 (0.020)	2.695 (0.012)
$\lambda = 0.75$	1.651 (0.017)	2.802 (0.012)

Notes: Each cell reports the mean IV estimate across five simulation seeds, with the seed-to-seed standard deviation in parentheses. Each seed uses 200,000 baseline firms, each copied into four equal-sized policy groups with common initial states and random shocks. The control group receives no change. In (a) the three treated groups receive cost increases $m_\kappa + \{0.00025, 0.00050, 0.00075\}$. In (b) they receive reductions in η of $m_\eta + \{0.0125, 0.0250, 0.0375\}$. (c) combines the middle cases from each panel; recurring-cost increases are multiplied by $1 - \lambda$ and η reductions by λ. $\widehat{\sigma}_h$ uses the response h periods after the policy change.

firms' adjustment to the new desired level of intangible capital after expanding to produce abroad, whereas the longer-term estimate $\widehat{\sigma}_{20}$ does. Thus, given enough time for the economy to adjust, the IV estimation in this experiment successfully recovers the structural parameter value. The size of the shock does not appear to play a material role; the estimates are very similar regardless of m_k.

5.7.3 Variable cost reductions

Panel (b) shows our second set of experiments, where the shock is to the variable production wedge, η. Firms in the treated groups experience a decrease in this parameter of $m_\eta + \{0.0125, 0.0250, 0.0375\}$ for $m_\eta \in \{0.05, 0.075, 0.10\}$. Here, the differences between the short-run and long-run estimates are magnified. The short-run estimate $\widehat{\sigma}_1$ falls to 1.6–1.7, and the long-run estimate $\widehat{\sigma}_{20}$ rises to 2.8–3.0, well above the true structural parameter value. The size of the shock matters more than in the first set of experiments: larger reductions in η yield larger estimates.

Why is the short-run elasticity estimate lower and the long-run estimate higher in this set of experiments? The reason is that reducing η has two separate effects on intangible intensity. First, it makes firms more likely to expand abroad, and the firms that do expand demand more intangible capital to exploit nonrivalry when $\sigma > 1$. This is the channel the IV estimation is designed to capture. Second, it increases demand for local tangible capital in the foreign affiliate, which in turn increases the marginal product of intangible capital. Importantly, this second channel operates even for firms that maintain the same number of foreign affiliates before and after the shock. In the short run, a firm that expands abroad in response to the η shock demands more tangible capital even though its stock of intangible capital has not yet adjusted, which pushes the intensity ratio downward. In the long run, its tangible capital demand grows along with its intangible capital, reinforcing the nonrivalry effect on the latter and leading to a larger long-run adjustment than when the shock is to $\bar{\kappa}_c$, pushing the intensity ratio upward.

5.7.4 Combined cost reductions

It is likely that the natural experiment in our empirical setting constitutes a shock to both the fixed and variable costs of multinational production. Our last set of experiments, shown in panel (c) of Table 5, explores how such a shock would play out in our model. shows the last set of experiments, where we combine the middle versions from the first two sets. Firms in the treated groups now experience both an increase in $\bar{\kappa}_c$ and a decrease in η, with a weight $\lambda \in \{0.25, 0.5, 0.75\}$ on the latter and the corresponding weight $1 - \lambda$ on the former. The short-run elasticity estimate $\widehat{\sigma}_1$ ranges from 1.65–1.8 and the long-run estimate $\widehat{\sigma}_{20}$ ranges from 2.56–2.8. These results confirm that our structural elasticity of 2.25 is compatible with reduced-form IV estimates in the range we reported in our empirical analysis in section 3.

6 Macroeconomic Implications

Both estimates place the elasticity of substitution between intangible capital and rival inputs above one. This section draws out two implications. First, growth accounting with common output elasticities attributes part of intangible deepening to productivity, even when intangible capital is measured and included as an input. Second, intangible deepening lowers the labor share while raising domestic employment.

We quantify these implications through an unanticipated, permanent 10% reduction in the purchase price of intangible investment, P. We compare the estimated technology, $\sigma = 2.25$, with a Cobb–Douglas benchmark. For the benchmark, we adjust μ to match the estimated economy's

baseline aggregate domestic labor share, holding the remaining parameters and production normalizations fixed. Each technology has its own baseline population. Within each specification, we compare the shocked economy with a control economy in which P remains unchanged, starting from the same firms and using common productivity, survival, entrant-characteristic, and affiliate-choice draws. Wages, tangible rental costs, and foreign operating conditions remain fixed. All reported responses compare the two paths five years after announcement. We use a price shock rather than a change in FDI barriers because it isolates the substitution channel: every firm faces the same change in relative input prices, and the responses do not depend on which firms expand abroad. Intangible deepening induced by multinational expansion, as in the experiments of Section 5.7, works through the same plant-level mechanism.

6.1 Aggregate Productivity Measurement

Intangible capital creates several difficulties for productivity measurement. Corrado et al. (2009) show that capitalizing intangible investment changes both measured output and the contribution of capital deepening. Crouzet and Eberly (2021) show how omitted intangibles and markups jointly distort measured productivity growth. Multinational production adds the problem of attributing output across countries: Guvenen, Mataloni Jr, Rassier, and Ruhl (2022) document how offshore profit shifting affects U.S. output and productivity measures. The scope of intangible use also matters for measuring capital services, as shown by Crouzet et al. (2026).

When $\sigma > 1$, the output elasticity of intangible capital differs across plants and changes as plants deepen. Growth accounting that imposes common, fixed output elasticities can then misattribute input contributions even when all inputs are observed. We compare such a measure with one based on the production function.

6.1.1 Nonrivalry and Output Elasticities

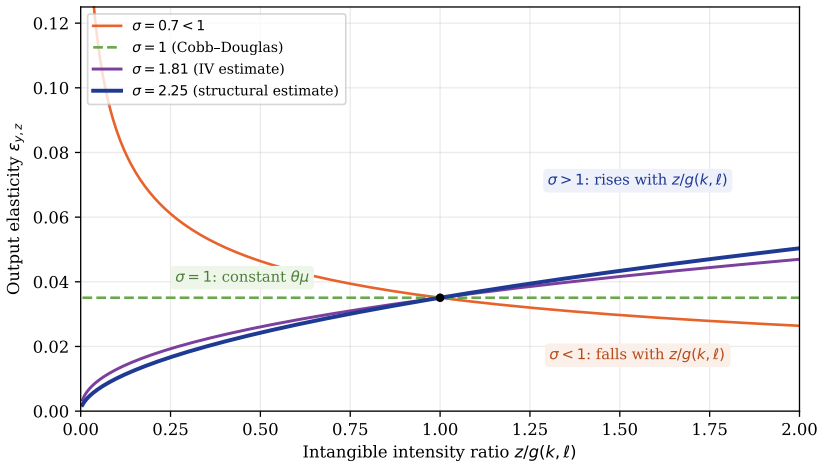
For a plant with output $y = f(a, z, k, \ell)$ given by (10), the output elasticity of intangible capital is

$$\varepsilon_{y,z} \equiv \frac{\partial \log f(a, z, k, \ell)}{\partial \log z} = \theta \frac{\mu^{1/\sigma} z^{(\sigma-1)/\sigma}}{\mu^{1/\sigma} z^{(\sigma-1)/\sigma} + (1-\mu)^{1/\sigma} g(k, \ell)^{(\sigma-1)/\sigma}}. \tag{31}$$

Since $g(k, \ell) = k^{\alpha} \ell^{1-\alpha}$, the output elasticities of tangible capital and labor are

$$\varepsilon_{y,k} = \alpha\big(\theta - \varepsilon_{y,z}\big), \qquad \varepsilon_{y,\ell} = (1-\alpha)\big(\theta - \varepsilon_{y,z}\big). \tag{32}$$

Figure 5: Intangible Intensity and the Output Elasticity of Intangible Capital



The three elasticities sum to θ. Under Cobb–Douglas production, $\sigma = 1$, the intangible elasticity is $\theta\mu$, and all three elasticities are constant across plants. When $\sigma > 1$, a higher ratio $z/g(k, \ell)$ raises the intangible elasticity and lowers the elasticities of tangible capital and labor. Figure 5 illustrates this relationship.

Nonrivalry makes the firm's intangible stock available at each of its production locations, while rival inputs are chosen locally. Differences in intangible intensity therefore imply differences in input contributions across plants. A common intangible elasticity understates the contribution of intangible deepening at plants with higher elasticities and overstates it at plants with lower elasticities, to first order. These errors can arise among firms with unchanged multinational scope. For finite changes, the elasticities also change as plants deepen.

6.1.2 Two Measures of Productivity Growth

We construct two measures using the same plant-level outputs and inputs. Let i index domestic plants and Δ denote the change between two dates. The production-function-consistent measure is

$$G_i^{\mathrm{PF}} = \Delta \log y_i - \Delta \log f(1, z_i, k_i, \ell_i). \tag{33}$$

Since $f(a, z, k, \ell) = a f(1, z, k, \ell)$, this measure recovers $\Delta \log a_i$. Evaluating the production function at both dates accounts for changes in output elasticities as the input mix changes.

The common-elasticity measure is

$$G_i^{\mathrm{CE}} = \Delta \log y_i - \bar{\varepsilon}_{y,z} \Delta \log z_i - \bar{\varepsilon}_{y,k} \Delta \log k_i - \bar{\varepsilon}_{y,\ell} \Delta \log \ell_i. \tag{34}$$

For the cohort used in the comparison, let $s_{i,0}$ be each plant's share of initial domestic output, normalized to sum to one. Set $\bar{\varepsilon}_{y,m} = \sum_i s_{i,0} \varepsilon_{y,m,i,0}$ for $m \in \{z, k, \ell\}$. The common elasticities sum to θ, so both measures include intangible capital and impose the same returns to scale. Capital inputs are measured in quantities; the decline in P does not mechanically change the measured stock z.

We aggregate using the same initial weights, $G^q = \sum_i s_{i,0} G_i^q$ for $q \in \{\mathrm{PF}, \mathrm{CE}\}$. Holding these weights fixed isolates differences in measured growth within plants. To first order, the measurement gap is

$$G^{\mathrm{CE}} - G^{\mathrm{PF}} \simeq \sum_i s_{i,0} \big(\varepsilon_{y,z,i,0} - \bar{\varepsilon}_{y,z}\big) \Delta \log \left(\frac{z_i}{g(k_i, \ell_i)} \right). \tag{35}$$

This expression is the weighted covariance between the initial intangible elasticity and subsequent intangible deepening. The gap is positive to first order when deepening is concentrated at plants with higher elasticities. For finite changes, changes in the elasticities also matter; the calculations below use the exact measures. Proportional growth in all inputs generates no gap, and under Cobb–Douglas the measures coincide for any change in inputs.

6.1.3 Quantitative Implications

We apply both measures to the intangible-price experiment described above. For each technology, we select the same cohort of initial producing domestic firms that survives through year five in the shocked and control simulations. The common elasticities and aggregation weights are computed once from this cohort's initial observations and used in both paths. Let superscripts S and C denote the shocked and control economies. Table 6 reports $100(G^{q,\mathrm{S}} - G^{q,\mathrm{C}})$ for each measure and their difference, in log percentage points of cumulative five-year growth.

Under the estimated technology, common elasticities generate an additional 0.46 log percentage points of measured productivity growth over five years, about 0.1 percentage point per year. The production-function-consistent response is zero: primitive productivity follows the same path in both economies, and the aggregation weights are fixed. The positive gap therefore reflects input contributions that the common-elasticity measure fails to remove. Cheaper intangible investment raises measured productivity even though the shock does not change primitive productivity. Under Cobb–Douglas, the common elasticities coincide with the production elasticities and the gap is

Table 6: Productivity Measurement Following Cheaper Intangible Investment

	Production function	Common elasticities	Gap
Estimated CES ($\sigma = 2.25$)	0.00	0.46	0.46
Cobb–Douglas ($\sigma = 1$)	0.00	0.00	0.00

Notes: Responses to an unanticipated, permanent 10% reduction in P, relative to a matched control path five years after announcement. Entries are cumulative growth differences in log percentage points, not annualized. The first column is the production-function-consistent measure and the second the common-elasticity measure; the gap is the second minus the first. Both measures use the same surviving initial producing firms and fixed initial output weights. Common elasticities are initial output-weighted means. The Cobb–Douglas benchmark matches the estimated economy's baseline aggregate domestic labor share. Results pool two simulations of 50,000 firm slots each; the CES gap ranges from 0.43 to 0.50 log percentage points across the two simulations.

zero.

This exercise isolates measurement within surviving plants. Applying the argument to BLS productivity statistics also requires accounting for the construction of capital services and their weights, and the attribution of output across countries.[15] The reported gap is a model accounting comparison, not an estimate of the correction to U.S. productivity statistics.

6.2 The Labor Share and Domestic Employment

Research on the decline in the labor share emphasizes technological change, globalization, and changes in market power, as reviewed by Grossman and Oberfield (2022). Karabarbounis and Neiman (2014) emphasize substitution toward capital following a decline in the relative price of investment. Autor et al. (2020) document the contribution of reallocation toward firms with lower labor shares. Our estimated elasticity concerns substitution between intangible capital and the joint bundle of tangible capital and labor. We examine how cheaper intangible investment affects the labor share and domestic employment through changes within plants and changes in the composition of domestic production.

6.2.1 Intangible Intensity and the Plant Labor Share

Let i index producing domestic plants. With competitive labor markets and the domestic output price normalized to one, the first-order condition for labor implies

$$s_{\ell,i} \equiv \frac{W\ell_i}{y_i} = \varepsilon_{y,\ell,i} = (1-\alpha)\big(\theta - \varepsilon_{y,z,i}\big). \tag{36}$$

[15] BLS capital input includes software, research and development, and artistic originals. Its growth accounts use changing cost-share weights in chained Törnqvist indexes (U.S. Bureau of Labor Statistics, 2026).

Its response to intangible intensity satisfies

$$\frac{\partial \log s_{\ell,i}}{\partial \log[z_i/g(k_i,\ell_i)]} = -\frac{\sigma-1}{\sigma}\frac{\varepsilon_{y,z,i}}{\theta}. \tag{37}$$

When $\sigma > 1$, intangible deepening lowers the plant's labor share. When $\sigma < 1$, it raises the labor share. Under Cobb–Douglas, every producing plant has the constant share $(1-\alpha)\theta(1-\mu)$.

Tangible capital's rental share, $Rk_i/y_i = \alpha(\theta - \varepsilon_{y,z,i})$, also falls with intangible intensity when $\sigma > 1$. Payments to labor relative to tangible capital remain fixed at $(1-\alpha)/\alpha$. The share of revenue remaining after payments to rival inputs is

$$\frac{y_i - W\ell_i - Rk_i}{y_i} = 1 - \theta + \varepsilon_{y,z,i}. \tag{38}$$

This surplus is measured before accounting for intangible capital and the costs of multinational operations. Its increase cannot be interpreted directly as an increase in pure profits. The distinction between capital costs and profits is central to Barkai (2020).

6.2.2 Intangible Capital and Domestic Employment

A lower labor share does not require lower employment. Let ℓ_i^* denote optimal domestic employment at given productivity and factor prices. The optimal tangible-capital–labor ratio is fixed at those prices, so employment moves with the rival-input bundle. The cross-partial derivative of output with respect to intangible capital and that bundle has the sign of $\theta - (\sigma-1)/\sigma$. Strict concavity in the rival bundle implies

$$\frac{\partial \ell_i^*}{\partial z_i} > 0 \quad \Longleftrightarrow \quad \sigma < \frac{1}{1-\theta}. \tag{39}$$

With $\theta = 0.7$, the threshold is approximately 3.33. The structural estimate, $\sigma = 2.25$, lies between one and this threshold. So does the IV estimate, and so does the upper end of a two-standard-error interval around the structural estimate, 2.89. An increase in intangible capital therefore raises domestic output and employment while lowering the plant's labor share: output grows faster than payroll. This combination depends on both the substitution elasticity and returns to scale. Under Cobb–Douglas, intangible accumulation also raises output and employment, but their proportional responses are equal at each plant.

Nonrivalry allows additional intangible capital to be used at home and abroad, linking invest-

ment across locations as in McGrattan and Prescott (2009). This provides a possible mechanism for the complementarity between foreign and domestic investment documented by Desai, Foley, and Hines (2009). Our experiment changes the cost of acquiring that common stock while holding foreign operating conditions fixed.

6.2.3 Aggregation

Define domestic output and employment as $Y^H = \sum_i y_i$ and $L^H = \sum_i \ell_i$, and let $\omega_i = y_i/Y^H$. The domestic labor share is

$$S_\ell \equiv \frac{WL^H}{Y^H} = \sum_i \omega_i s_{\ell,i} = (1-\alpha)\left(\theta - \sum_i \omega_i \varepsilon_{y,z,i}\right). \tag{40}$$

It can fall through intangible deepening within plants or through reallocation toward plants with lower labor shares. The latter is the accounting mechanism emphasized by Autor et al. (2020); here, differences in intangible intensity generate differences in plant labor shares with competitive pricing.

We compare the shocked and control economies at year five, using the current domestic population and including replacement firms once they begin production. Each producing plant has a counterpart in the other simulation, but its inputs, output, and labor share can differ. Let $s^{\mathrm{S}}_{\ell,i}$ and $s^{\mathrm{C}}_{\ell,i}$ denote these labor shares and ω^{S}_i and ω^{C}_i the corresponding domestic output shares. Define

$$\Delta s_{\ell,i} \equiv s^{\mathrm{S}}_{\ell,i} - s^{\mathrm{C}}_{\ell,i}, \qquad \bar{s}_{\ell,i} \equiv \frac{s^{\mathrm{S}}_{\ell,i} + s^{\mathrm{C}}_{\ell,i}}{2},$$
$$\Delta \omega_i \equiv \omega^{\mathrm{S}}_i - \omega^{\mathrm{C}}_i, \qquad \bar{\omega}_i \equiv \frac{\omega^{\mathrm{S}}_i + \omega^{\mathrm{C}}_i}{2}.$$

The aggregate difference, $\Delta S_\ell = S^{\mathrm{S}}_\ell - S^{\mathrm{C}}_\ell$, has the exact decomposition

$$\Delta S_\ell = \underbrace{\sum_i \bar{\omega}_i \Delta s_{\ell,i}}_{\text{Changes within plants}} + \underbrace{\sum_i \bar{s}_{\ell,i} \Delta \omega_i}_{\text{Changes in output shares}}. \tag{41}$$

The first term averages changes in plant labor shares; the second captures reallocation of domestic output. The averages across the two economies make the decomposition exact for finite changes. Under Cobb–Douglas, both terms are zero because all producing plants have the same labor share.

Table 7: Labor Share and Domestic Activity Following Cheaper Intangible Investment

Panel A: Labor-share response (percentage points)			
	Within plants	Output shares	Total
Estimated CES ($\sigma = 2.25$)	-0.81	-0.02	-0.83
Cobb–Douglas ($\sigma = 1$)	0.00	0.00	0.00
Panel B: Stock and domestic activity (percent)			
	Intangibles	Domestic output	Domestic employment
Estimated CES ($\sigma = 2.25$)	27.84	2.62	1.19
Cobb–Douglas ($\sigma = 1$)	13.39	2.41	2.41

Notes: Responses to an unanticipated, permanent 10% reduction in P, relative to a matched control path five years after announcement. Panel A reports the terms in (41), multiplied by 100. Panel B reports $100(Z_5^{\mathrm{S}}/Z_5^{\mathrm{C}} - 1)$, $100(Y_5^{H,\mathrm{S}}/Y_5^{H,\mathrm{C}} - 1)$, and $100(L_5^{H,\mathrm{S}}/L_5^{H,\mathrm{C}} - 1)$. Aggregates cover the full current population, including replacement entrants. Setup stocks count in Z; newborn firms first produce in their next surviving period. Wages, tangible rental costs, and foreign operating conditions are fixed. The Cobb–Douglas benchmark matches the estimated economy's baseline aggregate domestic labor share. Results pool two simulations of 50,000 firm slots each.

6.2.4 Quantitative Implications

Table 7 reports responses to the same 10% reduction in P. Panel A decomposes the change in the domestic labor share. Panel B reports changes in intangible capital, domestic output, and domestic employment. The intangible aggregate, Z, counts the stock once per firm, including setup stocks held by new entrants. Output and payroll include replacement firms once they begin production.[16]

Under the estimated technology, the labor share falls by 0.83 percentage points. Changes within plants account for about 98% of the decline; reallocation contributes only 0.02 percentage points. In this experiment, the aggregate decline therefore comes primarily from changes in plants' input mix, with little contribution from changes in their output shares.

The comparison with Cobb–Douglas shows how substitutability changes the composition of the response. Intangible capital rises by 27.84% under the estimated technology, compared with 13.39% under Cobb–Douglas. Domestic output rises by similar amounts, 2.62% and 2.41%, respectively. Domestic employment, however, rises by only 1.19% under the estimated technology, roughly half the 2.41% response under Cobb–Douglas. Cheaper intangible investment thus generates a similar percentage increase in output with a smaller percentage increase in employment.

[16]The labor share covers payroll and output in the model's domestic production. A comparison with national accounts requires consistent treatment of intangible investment and the labor used to produce it. Koh, Santaeulàlia-Llopis, and Zheng (2020) show how capitalizing intellectual property products affects the measured labor share.

Since wages are fixed across the two paths,

$$\frac{S_\ell^{\mathrm{S}}}{S_\ell^{\mathrm{C}}} = \frac{L^{H,\mathrm{S}}/L^{H,\mathrm{C}}}{Y^{H,\mathrm{S}}/Y^{H,\mathrm{C}}}.$$

The decline in the labor share reflects output growing faster than payroll, while employment remains above its control path. Under Cobb–Douglas, output and payroll rise in the same proportion and the labor share is unchanged. The estimated elasticity therefore matters for the employment and factor-income consequences of cheaper intangible investment even when the output responses are similar. Two caveats apply. Wages and rental rates are fixed; in general equilibrium, the increase in labor demand would raise wages, so the reported decline in the labor share is an upper bound for a given labor supply elasticity. And the shock is common to all firms, which is why reallocation contributes so little; a shock that falls differentially on intangible-intensive firms, such as a reduction in FDI barriers, would raise the contribution of the reallocation channel of Autor et al. (2020).

7 Conclusion

Intangible capital is paid for once and used at every location where a firm produces. This paper has drawn out three consequences of that fact.

First, non-rivalry turns the geographic scope of a firm into a shifter of the relative price of intangible capital, and the response of intangible intensity to scope identifies the elasticity of substitution between intangibles and rival inputs. Multinationals are more intangible intensive than domestic firms, the gap grows with the number of countries in which a firm produces, and intensity rises after firms first expand abroad. Using reductions in FDI barriers as instruments for scope, we estimate an elasticity of 1.81. Intangible capital and rival inputs are gross substitutes, and the rise of intangibles and the spread of multinational production reinforce each other.

Second, the structural estimate of 2.25, obtained from a dynamic model of heterogeneous firms and identified from the cross-section rather than from policy-driven changes in scope, is consistent with the IV estimate once adjustment is taken into account. Intangible capital is built gradually, so the within-firm response to a change in scope over a few years understates the long-run elasticity. Running the IV design on data simulated from the model reproduces the empirical estimates in the short run, and shows that the long-run bias depends on whether a policy changes the fixed or the variable cost of foreign production.

Third, an elasticity above one changes how intangible deepening appears in aggregate data.

Growth accounting with common output elasticities records productivity gains that are not there, even when intangible capital is fully measured, and cheaper intangibles lower the labor share within plants while raising employment. Neither effect exists under Cobb–Douglas, the assumption that quantitative models of intangible capital and multinational production typically make.

References

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2023): "A Theory of Falling Growth and Rising Rents," The Review of Economic Studies, 90, 2675–2702.

Alessandria, G. and H. Choi (2007): "Do Sunk Costs of Exporting Matter for Net Export Dynamics?" Quarterly Journal of Economics, 122, 289–336.

Altonji, J. G. and L. M. Segal (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," Journal of Business Economic Statistics, 14, 353–366.

Alviarez, V., J. Cravino, and N. Ramondo (2023): "Firm-Embedded Productivity and Cross-Country Income Differences," Journal of Political Economy, 131, 2289–2327.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," The Quarterly Journal of Economics, 132, 1553–1592.

Antrás, P. and S. R. Yeaple (2014): "Multinational Firms and the Structure of International Trade," in Handbook of International Economics, ed. by G. Gopinath, E. Helpman, and K. Rogoff, Elsevier, vol. 4 of Handbook of International Economics, chap. 0, 55–130.

Arkolakis, C., N. Ramondo, A. Rodríguez-Clare, and S. Yeaple (2018): "Innovation and Production in the Global Economy," American Economic Review, 108, 2128–2173.

Aum, S. and Y. Shin (2024): "Is Software Eating the World?" Working Paper 32591, National Bureau of Economic Research.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The Fall of the Labor Share and the Rise of Superstar Firms," The Quarterly Journal of Economics, 135, 645–709.

Barkai, S. (2020): "Declining Labor and Capital Shares," The Journal of Finance, 75, 2421–2463.

Benhabib, J., J. Perla, and C. Tonetti (2021): "Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier," Econometrica, 89, 2261–2301.

Bilir, L. K. and E. Morales (2020): "Innovation in the Global Firm," Journal of Political Economy, 128, 1566–1625.

Corrado, C., C. Hulten, and D. Sichel (2009): "Intangible Capital and U.S. Economic Growth," Review of Income and Wealth, 55, 661–685.

Crouzet, N. and J. Eberly (2021): "Intangibles, Markups, and the Measurement of Productivity Growth," Journal of Monetary Economics, 124, S92–S109.

Crouzet, N., J. Eberly, A. Eisfeldt, and D. Papanikolaou (2026): "Intangible Capital, Firm Scope, and Growth," Working paper.

Crouzet, N., J. C. Eberly, A. L. Eisfeldt, and D. Papanikolaou (2022): "The Economics of Intangible Capital," Journal of Economic Perspectives, 36, 29–52.

Das, S., M. J. Roberts, and J. R. Tybout (2007): "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," Econometrica, 75, 837–873.

Desai, M. A., C. F. Foley, and J. R. Hines (2009): "Domestic Effects of the Foreign Activities of US Multinationals," American Economic Journal: Economic Policy, 1, 181–203.

Duarte, V. and J. Fonseca (2026): "AI for Structural Estimation," Working Paper 35283, National Bureau of Economic Research.

Dyrda, S., G. Hong, and J. B. Steinberg (2026): "Endogenous Firm Scope via Organizational Complexity," Working paper.

Eisfeldt, A. L. and D. Papanikolaou (2013): "Organization Capital and the Cross-Section of Expected Returns," The Journal of Finance, 68, 1365–1406.

Grossman, G. M. and E. Oberfield (2022): "The Elusive Explanation for the Declining Labor Share," Annual Review of Economics, 14, 93–124.

Guvenen, F., R. J. Mataloni Jr, D. G. Rassier, and K. J. Ruhl (2022): "Offshore Profit Shifting and Aggregate Measurement: Balance of Payments, Foreign Investment, Productivity, and the Labor Share," American Economic Review, 112, 1848–1884.

Helpman, E., M. J. Melitz, and S. R. Yeaple (2004): "Export versus FDI with heterogeneous firms," American Economic Review, 94, 300–316.

Jerbashian, V. (2024): "On the Elasticity of Substitution between Labor and ICT and IP Capital and Traditional Capital," GLO Discussion Paper Series 1523, Global Labor Organization (GLO).

Jones, C. I. and C. Tonetti (2020): "Nonrivalry and the Economics of Data," American Economic Review, 110, 2819–2858.

Karabarbounis, L. and B. Neiman (2014): "The Global Decline of the Labor Share," The Quarterly Journal of Economics, 129, 61–103.

Kleinman, B. (2025): "Wage inequality and the spatial expansion of firms," Tech. rep., Tech. rep., mimeo.

Klenow, P. J. and H. Li (2025): "Entry Costs Rise with Growth," Journal of Political Economy Macroeconomics, 3, 43–74.

Koh, D., R. Santaeulàlia-Llopis, and Y. Zheng (2020): "Labor Share Decline and Intellectual Property Products Capital," Econometrica, 88, 2609–2628.

Lashkari, D., A. Bauer, and J. Boussard (2024): "Information Technology and Returns to Scale," American Economic Review, 114, 1769–1815.

Li, W. C. Y. and B. H. Hall (2016): "Depreciation of Business R&D Capital," Working Paper 22473, National Bureau of Economic Research.

McGrattan, E. R. and E. C. Prescott (2009): "Openness, technology capital, and development," Journal of Economic Theory, 144, 2454–2476.

——— (2010): "Technology Capital and the US Current Account," American Economic Review, 100, 1493–1522.

Oberfield, E. and D. Raval (2021): "Micro Data and Macro Technology," Econometrica, 89, 703–732.

Oberfield, E., E. Rossi-Hansberg, P.-D. Sarte, and N. Trachter (2024): "Plants in space," Journal of Political Economy, 132, 867–909.

O'Mahony, M., C. Corrado, J. Haskel, M. Iommi, C. Jona-Lasinio, and M. Mas (2018): "Advancements in measuring intangibles for European economies," EURONA, 2017.

Perla, J., C. Tonetti, and M. E. Waugh (2021): "Equilibrium Technology Diffusion, Trade, and Growth," American Economic Review, 111, 73–128.

Peters, R. H. and L. A. Taylor (2017): "Intangible capital and the investment-q relation," Journal of Financial Economics, 123, 251–272.

Ramondo, N. and A. Rodríguez-Clare (2013): "Trade, Multinational Production, and the Gains from Openness," Journal of Political Economy, 121, 273–322.

Ramondo, N., A. Rodríguez-Clare, and F. Tintelnot (2015): "Multinational Production: Data and Stylized Facts," American Economic Review, 105, 530–536.

Ruhl, K. J. and J. L. Willis (2017): "NEW EXPORTER DYNAMICS," International Economic Review, 58, 703–725.

Spencer, A. H. (2021): "Policy Effects of International Taxation on Firm Dynamics and Capital Structure," The Review of Economic Studies, 89, 2149–2200.

——— (2022): "Policy Effects of International Taxation on Firm Dynamics and Capital Structure," The Review of Economic Studies, 89, 2149–2200.

Steinberg, J. B. (2023): "Export market penetration dynamics," Journal of International Economics, 145, 103807.

Sun, L. and S. Abraham (2021): "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," Journal of Econometrics, 225, 175–199, themed Issue: Treatment Effect 1.

Tintelnot, F. (2017): "Global Production with Export Platforms*," The Quarterly Journal of Economics, 132, 157–209.

U.S. Bureau of Labor Statistics (2026): "Total Factor Productivity: Technical Notes," Updated March 19, 2026.

Appendix

A Theoretical Results and Proofs

This appendix proves Proposition 1 and collects the comparative-statics results of the static theory in Section 2, together with their proofs. Throughout, we maintain Assumption 1. For the comparative-statics results we impose two additional conditions, which Proposition 1 does not require: the inputs are gross substitutes, $\sigma > 1$, consistent with our estimates in Section 3, and the weight on the rival input, $1 - \mu$, is sufficiently large.

For notational convenience, we introduce the following definitions.

Definition 1. *We define the elasticity terms as follows:*

1. *The elasticity of an endogenous variable y with respect to a variable c is defined as $\varepsilon_{y,c} = \frac{\partial \ln y}{\partial \ln c}$.*
2. *The combined elasticity of an endogenous variable y with respect to variables c_1 and c_2 is defined as*

$$\varepsilon_{y,(c_1,c_2)} = \varepsilon_{y,c_1} + \varepsilon_{y,c_2} + \frac{\partial^2 \ln y}{\partial \ln c_1 \partial \ln c_2}.$$

A.1 Statements

The first set of results concerns reductions in the price of intangible capital P and the expansion cost $\bar{\kappa}$.

Proposition 2. *When the price of intangible capital P or the expansion cost $\bar{\kappa}$ falls, the following statements hold under the assumptions:*

1. *all firms increase in intangible intensity, i.e. $\varepsilon_{s^z,c} < 0$, $c \in \{P, \bar{\kappa}\}$;*
2. *all firms expand into more foreign countries, i.e. $\varepsilon_{n,c} < 0$, $c \in \{P, \bar{\kappa}\}$.*

The following lemma illustrates the role of multinational expansion in mediating the impact on intangible investment.

Lemma 2. *We have $\frac{\partial s^z}{\partial c} = \frac{\partial s^z}{\partial c}|_n + \frac{\partial s^z}{\partial n}\frac{\partial n}{\partial c}$, $c \in \{P, \bar{\kappa}\}$. Endogenous multinational expansion amplifies the effect of cost reductions on the intangible intensity of a firm, i.e. $\frac{\partial s^z}{\partial n}\frac{\partial n}{\partial c} < 0$.*

The next proposition formalizes the complementary effects of reductions in P and $\bar{\kappa}$ on s^z and n.

Proposition 3. *For $y = \{s^z, n\}$, reductions in $\bar{\kappa}$ and P are complementary: their joint effect on y is super-additive,*

$$|\varepsilon_{y,(\bar{\kappa},P)}| > |\varepsilon_{y,\bar{\kappa}}| + |\varepsilon_{y,P}|.$$

The inequality is expressed in absolute values because the elasticities are negative: reductions in the costs increase s^z and n. The complementarity stems from the interdependence of multinational expansion and intangible investment: cheaper intangibles raise per-location profits and amplify the expansion response, while cheaper expansion raises the number of locations over which the non-rival input is deployed and amplifies the investment response.

The last set of results concerns the role of firm heterogeneity.

Proposition 4. *The following statements are true under the assumptions:*

1. *Intangible intensity increases with firm productivity, i.e. $\varepsilon_{s^z,a} > 0$.*
2. *The global scope of production increases with firm productivity, i.e. $\varepsilon_{n,a} > 0$.*

Proposition 5. *The following statements are true under the assumptions for $y = \{s^z, n\}$:*

1. *A decrease in P increases n and s^z by more for high-productivity firms, i.e.*

$$\left|\varepsilon_{y,(a,P)}\right| > |\varepsilon_{y,a}| + |\varepsilon_{y,P}|\,.$$

2. *A decrease in $\bar{\kappa}$ increases n and s^z by more for high-productivity firms, i.e.*

$$\left|\varepsilon_{y,(a,\bar{\kappa})}\right| > |\varepsilon_{y,a}| + |\varepsilon_{y,\bar{\kappa}}|\,.$$

The two propositions state that more productive firms choose higher intangible intensity and a larger global footprint, and that cost reductions have larger effects for them. The intuition is the same in both cases: the impact of a cost decline hinges on the firm's capacity to expand, which rises with productivity.

A.2 Proofs

Preliminary derivations. We first derive the two equilibrium conditions used repeatedly below. Let

$$\rho \equiv \frac{\sigma-1}{\sigma} < 1, \quad \rho \neq 0, \qquad \Xi(z,\ell) \equiv \mu^{1/\sigma} z^{\rho} + (1-\mu)^{1/\sigma} \ell^{\rho}, \qquad f(a,z,\ell) = a\,\Xi(z,\ell)^{\theta/\rho}.$$

The firm's problem is

$$\max_{\ell,z,n} \left\{ (1+n)\big[p\, f(a,z,\ell) - W\ell\big] - Pz - \bar{\kappa}\, n^{\gamma} \right\}.$$

The first-order conditions with respect to ℓ and z are

$$(1+n)\, p\, a\, \theta\, \Xi(z,\ell)^{\frac{\theta}{\rho}-1} (1-\mu)^{1/\sigma} \ell^{\rho-1} = (1+n)\, W, \tag{42}$$

$$(1+n)\, p\, a\, \theta\, \Xi(z,\ell)^{\frac{\theta}{\rho}-1} \mu^{1/\sigma} z^{\rho-1} = P. \tag{43}$$

The rival input is hired separately in each of the $1+n$ locations, so its marginal cost scales with $1+n$, whereas the non-rival input is purchased once and used everywhere. Dividing (43) by (42) and using $\rho - 1 = -1/\sigma$ yields

$$\frac{P}{(1+n)W} = \left(\frac{\mu}{1-\mu}\right)^{1/\sigma} \left(\frac{z}{\ell}\right)^{-1/\sigma}, \qquad \text{so} \qquad \frac{z}{\ell} = \frac{\mu}{1-\mu} \left(\frac{(1+n)W}{P}\right)^{\sigma}. \tag{44}$$

Because $s^z \equiv \frac{Pz}{(1+n)W\ell + Pz} = \left[1 + \frac{(1+n)W}{P}\frac{\ell}{z}\right]^{-1}$, substituting z/ℓ from (44) gives

$$s^z = \left[\frac{1-\mu}{\mu} \left(\frac{W}{P}\right)^{1-\sigma} (1+n)^{1-\sigma} + 1\right]^{-1}, \tag{45}$$

which is equation (4). By the envelope theorem applied to the choice of ℓ, the first-order condition with respect to n is

$$\tilde{\pi}(a,z) = \gamma \bar{\kappa}\, n^{\gamma-1}, \tag{46}$$

hence

$$n = \left(\frac{\tilde{\pi}(a,z)}{\gamma\bar{\kappa}}\right)^{\frac{1}{\gamma-1}}, \tag{47}$$

which is equation (3). Equations (45) and (47) imply three sign restrictions used throughout:

$$\frac{\partial s^z}{\partial n} > 0, \qquad \frac{\partial n}{\partial \tilde{\pi}} > 0, \qquad \frac{\partial \tilde{\pi}(a,z)}{\partial a} > 0. \tag{48}$$

The first inequality follows from $\sigma > 1$, so $1-\sigma < 0$ in (45); the second from $\gamma > 1$ in (47); and the third from $f(a,z,\ell) > 0$.

Proof of Proposition 1. Write (45) as $s^z = [\Lambda(1+n)^{1-\sigma}+1]^{-1}$ with $\Lambda \equiv \frac{1-\mu}{\mu}\left(\frac{W}{P}\right)^{1-\sigma} > 0$. Differentiating with respect to n gives

$$\frac{\partial s^z}{\partial n} = (\sigma-1)\,\Lambda(1+n)^{-\sigma}\left[\Lambda(1+n)^{1-\sigma}+1\right]^{-2},$$

whose sign is the sign of $\sigma - 1$. Hence s^z is strictly increasing in n if $\sigma > 1$ and strictly decreasing if $\sigma < 1$. For $\sigma = 1$, $(1+n)^{1-\sigma} = 1$ and n drops out of (45), so s^z is invariant to scope. The argument uses only Assumption 1. □

Proof of Proposition 2. We prove the two comparative statics for $c \in \{P, \bar{\kappa}\}$. Fix first n. From (44), a rise in P lowers z/ℓ, and because $1-\sigma < 0$, (45) shows that a rise in P lowers s^z at fixed n, so $\partial s^z/\partial P|_n < 0$. Likewise, because $1-\sigma < 0$, (45) implies $\partial s^z/\partial n > 0$. Next, (47) gives $\partial n/\partial \bar{\kappa} < 0$ immediately. Moreover, a rise in P raises the direct cost of intangible capital, reduces the optimal z, lowers per-location profit $\tilde{\pi}(a,z)$, and therefore lowers n through (47). Hence $\partial n/\partial P < 0$.

To complete the proof for s^z, apply the total derivative:

$$\frac{\partial s^z}{\partial c} = \left.\frac{\partial s^z}{\partial c}\right|_n + \frac{\partial s^z}{\partial n}\frac{\partial n}{\partial c}.$$

For $c = P$, both terms are negative. For $c = \bar{\kappa}$, the direct effect at fixed n is zero and the indirect term is negative because $\partial s^z/\partial n > 0$ and $\partial n/\partial \bar{\kappa} < 0$. Therefore $\varepsilon_{s^z,c} < 0$ and $\varepsilon_{n,c} < 0$ for $c \in \{P, \bar{\kappa}\}$. □

Proof of Lemma 2. The decomposition

$$\frac{\partial s^z}{\partial c} = \left.\frac{\partial s^z}{\partial c}\right|_n + \frac{\partial s^z}{\partial n}\frac{\partial n}{\partial c}$$

is the chain rule applied to $s^z = s^z(n,c)$. It remains to sign the amplification term. Equation (45) implies $\partial s^z/\partial n > 0$ because $\sigma > 1$. Proposition 2 established that $\partial n/\partial c < 0$ for $c \in \{P, \bar{\kappa}\}$.

Hence

$$\frac{\partial s^z}{\partial n}\frac{\partial n}{\partial c} < 0.$$

This is the claimed amplification effect: endogenous expansion magnifies the impact of lower costs on intangible intensity. □

Proof of Proposition 3. Let $y \in \{s^z, n\}$. By definition,

$$\varepsilon_{y,(\bar{\kappa},P)} = \varepsilon_{y,\bar{\kappa}} + \varepsilon_{y,P} + \frac{\partial^2 \ln y}{\partial \ln \bar{\kappa}\, \partial \ln P}.$$

It is therefore sufficient to show that the cross-partial term is negative. We first consider n. Taking logs in (47) gives

$$\ln n = \frac{1}{\gamma - 1}\big(\ln \tilde{\pi}(a,z) - \ln \gamma - \ln \bar{\kappa}\big).$$

Hence

$$\frac{\partial \ln n}{\partial \ln P} = \frac{1}{\gamma - 1}\frac{\partial \ln \tilde{\pi}}{\partial \ln P} < 0.$$

Because lower $\bar{\kappa}$ raises n, and $\partial s^z/\partial n > 0$, the induced rise in n increases the sensitivity of profits to P: when expansion is cheaper, the firm operates in more locations and the return to intangible capital is higher. Therefore

$$\frac{\partial^2 \ln \tilde{\pi}}{\partial \ln \bar{\kappa}\, \partial \ln P} > 0 \qquad \Longrightarrow \qquad \frac{\partial^2 \ln n}{\partial \ln \bar{\kappa}\, \partial \ln P} < 0.$$

Thus $|\varepsilon_{n,(\bar{\kappa},P)}| > |\varepsilon_{n,\bar{\kappa}}| + |\varepsilon_{n,P}|$.

Now consider s^z. Since s^z is increasing in n by (45), the same interaction carries over: a reduction in $\bar{\kappa}$ raises n, and a larger n strengthens the effect of lower P on the optimal intangible share. Therefore

$$\frac{\partial^2 \ln s^z}{\partial \ln \bar{\kappa}\, \partial \ln P} < 0,$$

which implies

$$|\varepsilon_{s^z,(\bar{\kappa},P)}| > |\varepsilon_{s^z,\bar{\kappa}}| + |\varepsilon_{s^z,P}|.$$

Combining the two cases proves the proposition. □

Proof of Proposition 4. Equation (47) shows that n is strictly increasing in per-location profit $\tilde{\pi}(a,z)$. Since

$$\tilde{\pi}(a,z) = \max_{\ell} \{p\, f(a,z,\ell) - W\ell\}$$

and $f(a, z, \ell) = a\,\Xi(z, \ell)^{\theta/\rho}$ is linear in a, the envelope theorem gives $\partial\tilde{\pi}/\partial a = p\,\Xi(z, \ell)^{\theta/\rho} > 0$. Therefore higher-productivity firms choose larger n, so $\varepsilon_{n,a} > 0$.

Because (45) implies $\partial s^z/\partial n > 0$, any increase in productivity that raises n also raises the optimal intangible intensity. Hence $\varepsilon_{s^z,a} > 0$. This proves both claims. □

Proof of Proposition 5. Let $y \in \{s^z, n\}$. We need to show that the effects of lower P and lower $\bar{\kappa}$ are larger for higher-productivity firms. Start with the interaction between a and P. Since $\partial\tilde{\pi}/\partial a > 0$, higher productivity raises the profit base from which the firm expands. Because n is increasing in profits, the effect of a lower P on n is stronger when a is higher. Therefore

$$\frac{\partial^2 \ln n}{\partial \ln a\, \partial \ln P} < 0.$$

Since s^z is increasing in n, the same argument implies

$$\frac{\partial^2 \ln s^z}{\partial \ln a\, \partial \ln P} < 0.$$

Hence

$$|\varepsilon_{y,(a,P)}| > |\varepsilon_{y,a}| + |\varepsilon_{y,P}|.$$

An analogous argument applies to the pair $(a, \bar{\kappa})$. A more productive firm has higher profits, so it operates on a steeper part of the expansion schedule (47). A reduction in $\bar{\kappa}$ therefore induces a larger increase in n for more productive firms, which in turn induces a larger rise in s^z. Thus

$$\frac{\partial^2 \ln n}{\partial \ln a\, \partial \ln \bar{\kappa}} < 0, \qquad \frac{\partial^2 \ln s^z}{\partial \ln a\, \partial \ln \bar{\kappa}} < 0,$$

and therefore

$$|\varepsilon_{y,(a,\bar{\kappa})}| > |\varepsilon_{y,a}| + |\varepsilon_{y,\bar{\kappa}}|.$$

This proves the proposition. □

B Data and Empirical Details

B.1 Data

Firm financials. Firm-level data are from Compustat North America Annual Fundamentals for 1995–2023. When a firm-year has both an industrial and a financial-services record, we keep the industrial record. We drop firm-years without a NAICS code or with non-positive net sales, and keep firm-years whose intangible intensity, defined in (49), lies between 0 and 1. Dollar amounts are in 2009 dollars. NAICS sectors 31–33, 44–45 and 48–49 are each treated as one 2-digit sector. The panel contains 190,357 firm-years of 19,971 firms.

Intangible capital. Balance-sheet intangibles are Compustat's total intangible assets, including goodwill. Knowledge capital is the stock constructed by Peters and Taylor (2017) and distributed by WRDS, which accumulates R&D expenditure by the perpetual-inventory method using the BEA's industry-specific R&D depreciation rates, and 15 percent for industries without one. We do not include organization capital. A firm-year has no measured intangible capital when it reports zero intangible assets and has no knowledge capital; this is the case for 17 percent of the panel, and 88 percent of these firm-years report no R&D.

Geographic scope. Exhibit 21 of Form 10-K lists a company's subsidiaries and each one's jurisdiction of incorporation. We use the subsidiary data WRDS extracts from these exhibits, date each filing by its filing year, and link filings to Compustat by gvkey, with a CIK–gvkey link table for filings without one. n_{it} is the number of distinct countries other than the United States in which firm i lists subsidiaries in year t. Firm-years without an Exhibit 21 record have $n_{it} = 0$.

B.2 Variables

Intangible intensity, equation (5) of the paper, is the share of intangible capital in total capital,

$$s_{it}^{z} = \frac{\text{INTAN}_{it} + G_{it}}{\text{INTAN}_{it} + G_{it} + \text{PPEGT}_{it}}, \tag{49}$$

where INTAN is balance-sheet intangibles, G is knowledge capital and PPEGT is gross property, plant and equipment. A firm-year is multinational if the firm lists at least one foreign subsidiary, $\text{MNE}_{it} = \mathbf{1}\{n_{it} > 0\}$. Scope enters equation (7) as $\log(1 + n_{it})$, which is zero for domestic firms, $\text{Post}_t = \mathbf{1}\{t \geq 2006\}$, and firm size is $\log(1 + y_{it})$, with y_{it} net sales in millions of 2009 dollars.

B.3 Samples

Table 1 uses the balanced panel of the 942 firms present in every year from 1995 to 2023 (27,318 firm-years; 11,925 MNE firm-years of 637 firms in Columns (4)–(6)), including firm-years without intangible capital. It compares average intensity across firms before and after 2005 without firm fixed effects, and the set of listed firms changes substantially over the period: 39 percent of firm-years before 2006 belong to firms that leave the panel by 2005, and 35 percent of firm-years after 2005 belong to firms that enter after 2005. Holding the set of firms fixed prevents entry and exit from appearing as a change in the multinational gap. In the full panel the gap in Column (3) is 7.9 percentage points and widens by 1.9 points after 2005, against 5.6 and 4.7 in the balanced panel.

Table 2 uses all firms but keeps only firm-years with $0.01 < s^z_{it} < 0.99$, since its dependent variable, $\log((1 - s^z_{it})/s^z_{it})$, is undefined at 0 and 1. This removes 38,959 firm-years with $s^z \leq 0.01$, of which 31,691 have no intangible capital, and 2,716 with $s^z \geq 0.99$, of which 1,391 have no PP&E, leaving 148,682 firm-years of 16,516 firms. The two estimation samples share 22,026 firm-years of 861 firms.

B.4 Summary statistics

Table 8 reports summary statistics for the Table 2 sample. Three features matter for the estimates. First, scope is skewed: 65 percent of firm-years are domestic, so the median of $\log(1 + n)$ is zero, and among MNE firm-years the median firm lists subsidiaries in 5 foreign countries and the 90th percentile in 28. Second, balance-sheet intangibles are mostly goodwill (55 percent on average from 2002, when Compustat separates it), and knowledge capital is zero in 38 percent of firm-years. Third, the panel is unbalanced: firms are observed for 9 years on average and 6 at the median.

Table 8: Summary Statistics: Estimation Sample of Table 2

	Mean	Std. dev.	P10	Median	P90
Panel A. Intangible intensity and scope					
Intangible intensity, s^z	0.483	0.307	0.059	0.496	0.891
Domestic firm-years	0.472	0.316	0.048	0.473	0.899
MNE firm-years	0.506	0.286	0.093	0.530	0.876
$\log\big((1-s^z)/s^z\big)$	0.16	1.85	-2.10	0.01	2.77
Foreign countries, n (MNE firm-years)	10.5	13.5	1.0	5.0	28.0
$\log(1+n)$	0.66	1.08	0.00	0.00	2.56
Panel B. Balance sheet (millions of 2009 dollars)					
Net sales	3,444.2	14,970.1	4.7	226.8	6,294.6
Gross property, plant and equipment	2,782.0	14,894.8	2.0	70.4	3,949.7
Balance-sheet intangibles	1,046.1	5,596.6	0.0	24.4	1,657.3
Knowledge capital	371.5	2,845.2	0.0	2.8	288.6
Goodwill share of balance-sheet intangibles[a]	0.546	0.359	0.000	0.634	0.993
Panel C. Shares of firm-years and panel structure					
Multinational (MNE)	0.346				
Positive knowledge capital	0.619				
Years per firm	9.0	7.7	2.0	6.0	22.0
Firm-years					148,682
Firms					16,516

Notes: The sample is the estimation sample of Table 2: firm-years with $0.01 < s^z < 0.99$, 1995–2023. Intangible intensity s^z is defined in (49). n is the number of distinct non-U.S. countries of incorporation among subsidiaries listed in Exhibit 21, summarized over MNE firm-years.

[a] Goodwill over balance-sheet intangibles, computed from 2002, when Compustat separates goodwill from other intangibles, over firm-years with positive balance-sheet intangibles.

B.5 Specification of Table 1

Columns (1)–(3) are OLS regressions on all firm-years of the balanced panel,

$$s^z_{it} = \beta_1 \,\mathrm{MNE}_{it} + \beta_2 \,\mathrm{MNE}_{it} \times \mathrm{Post}_t + \gamma \log(1+y_{it}) + \delta_t + \eta_{k(i)} + \varepsilon_{it}, \tag{50}$$

and columns (4)–(6) are OLS regressions on its MNE firm-years,

$$s_{it}^{z} = \beta_1 \log(1 + n_{it}) + \beta_2 \log(1 + n_{it}) \times \text{Post}_t + \gamma \log(1 + y_{it}) + \delta_t + \eta_{k(i)} + \varepsilon_{it}. \tag{51}$$

All columns include year fixed effects δ_t, which absorb the main effect of Post_t. Columns (2), (3), (5) and (6) control for firm size, and columns (3) and (6) add 2-digit NAICS fixed effects $\eta_{k(i)}$. Regressions are unweighted. Standard errors are clustered by year (29 clusters) in columns (1), (2), (4) and (5), and by 2-digit NAICS $\times$ year cell in columns (3) and (6).