Intangible Capital, Multinational Production, and Productivity Growth*

Sebastian Dyrda
University of Toronto

Guangbin Hong
Michigan State University

Muhammad Ali Sajid
University of Toronto

Joseph B. Steinberg
University of Toronto & NBER

March 26, 2026

Link to the current version

Abstract

Over the past three decades, intangible capital and multinational production have risen together and are increasingly concentrated among the most productive firms. We document these patterns using firm-level data and develop a quantitative general equilibrium model in which firms invest in nonrival intangible capital and establish foreign subsidiaries. Nonrivalry generates scalable complementarity: multinational expansion raises the return to intangible investment, while intangible intensity strengthens incentives to expand globally. Our calibration reveals that intangible and rival inputs are gross substitutes, with an elasticity of substitution of 1.8. This implies that when the price of intangible capital falls, firms substitute away from rival factors towards intangible capital, increasing intangible intensity at the firm level and raising output elasticities at the plant level. A structural accounting exercise attributes the joint rise in both trends primarily to declining intangible investment costs, with globalization frictions playing a secondary but amplifying role. A production-function-consistent measure of aggregate TFP reveals that these forces leave true productivity essentially unchanged. What standard measurement approaches attribute to efficiency gains is actually factor deepening in a nonrival input whose contribution is mismeasured—understated by the Solow residual and overstated by the BEA/BLS.

Keywords: Intangible capital; multinational production; total factor productivity; firm heterogeneity; globalization.

JEL Codes: E23, F23, O33, O41, H25, D24

*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.

1 Introduction

Over the past four decades, two key trends have dramatically reshaped the global economy: the rise of intangible capital and the expansion of multinational production. We argue that these trends are tightly linked through the nonrival nature of intangible capital and that their interaction has first-order implications for aggregate productivity measurement. Unlike tangible capital, intellectual property, software, and other intangible assets can be deployed simultaneously across multiple production locations. This scalability implies that multinational firms have stronger incentives to invest in intangible capital than purely domestic firms, while firms with higher intangible intensity have greater incentives to expand their production footprint across borders. As a result, reductions in barriers to multinational production raise intangible intensity, while intangible-biased technological change or declines in the cost of producing intangible capital increase multinational activity. Through these channels, the joint rise of intangibles and multinational production generates factor deepening in a nonrival input that standard productivity measurement misattributes as efficiency gains. We develop a quantitative model to discipline the importance of these forces, to measure their contribution to the observed rise in intangible intensity and multinational production, and to assess the implications for aggregate productivity measurement.

Our first contribution is to document these interconnected trends using linked firm-level data from Compustat and U.S. Securities and Exchange Commission (SEC) disclosures. Intangible intensity—measured as the ratio of intangible assets (including capitalized R&D) to total assets—more than doubled among U.S. public firms between 1995 and 2023, rising from 26% to 47%. This increase occurred broadly across industries and is driven primarily by within-industry changes rather than sectoral reallocation. Over the same period, multinational activity expanded sharply: the share of firms operating foreign subsidiaries rose from 7% to 30%, and the average foreign sales share of multinational enterprises (MNEs) increased from 10% to 22%.

Cross-sectional evidence reveals that these trends are closely linked at the firm level. Multinational firms exhibit substantially higher intangible intensity than purely domestic firms, with an average gap of 6 percentage points that doubles after 2005, and firms operating in a larger number of countries display even higher intangible intensity. These relationships are strongly structured by firm productivity: more productive firms are both more intangible-intensive and more likely to operate as multinationals. An event-study framework that traces within-firm changes around multinational entry confirms that the link is not purely cross-sectional: firms that become multinationals experience a sustained increase in intangible intensity of 6 percentage points following entry, consistent with multinational expansion raising the return to intangible investment

by increasing the scale over which intangibles can be deployed. Together, the cross-sectional and dynamic patterns indicate that the rise of intangible capital and multinational production is concentrated among high-productivity firms and reflects the joint role of productivity heterogeneity, scalable nonrival inputs, and endogenous organizational choices.

Motivated by these empirical patterns, we develop a structural framework that links the rise of intangible capital and multinational expansion within a unified economic environment. This framework, our second contribution, models firms' decisions to invest in intangibles and expand across borders as jointly determined outcomes of productivity heterogeneity, scalable technologies, and globalization frictions. By embedding these choices in a multi-region general equilibrium setting, the model captures equilibrium feedbacks, selection into multinational activity, and the interaction between firm organization and aggregate outcomes. A central feature of the framework is the nonrival nature of intangible capital: intangibles are developed at the firm's headquarters but can be deployed simultaneously across multiple production locations, implying that the return to intangible investment rises with the firm's multinational scope. Within this framework, reductions in the cost of producing intangible capital and declines in barriers to multinational expansion act as scalable complements: each force independently raises intangible intensity and the propensity to operate abroad, but their joint effect is superlinear.

Our third contribution is a quantitative analysis that disciplines the model and connects its firm-level mechanisms to aggregate trends. The analysis is designed to serve two purposes: to organize the data around a coherent structural framework, and to quantify the relative importance of technological change and globalization in driving the observed co-movement of intangible intensity and multinational activity. A key parameter is the elasticity of substitution between nonrival intangible capital and rival local inputs, σ, which governs how strongly firms shift toward intangibles as multinational scope expands. We identify σ in two complementary ways and obtain a robust finding that $\sigma > 1$: intangible and rival inputs are gross substitutes. In the model, we calibrate σ by targeting the 15-percentage-point gap in intangible intensity between multinational and non-multinational firms. We validate this approach by estimating directly in the firm-level data a structural relationship linking intangible intensity to the number of foreign subsidiaries, instrumenting the latter with exposure to the 2004 EU enlargement and the staggered expansion of bilateral investment treaty networks. Our model calibration and empirical estimation both yield estimates of 1.8, confirming gross substitutability. This finding has two consequences: it amplifies the scalable complementarity between intangible investment and multinational expansion, and it implies that firms with different multinational scope vary in their output elasticities—a dimension

of firm heterogeneity that, as we show, is central to aggregate productivity measurement.

We calibrate the model to match key features of U.S. publicly listed firms over the period 2005–2022 and conduct a structural accounting exercise to identify the forces behind the observed trends. Comparing pre- and post-2005 steady states, we back out changes in the cost of intangible investment and in globalization frictions that rationalize the joint rise in intangible intensity and multinational activity. The main finding is that declining intangible investment costs are the primary driver of both trends: the model attributes 8.9 of the 10.7 percentage point rise in the MNE share and nearly the entirety of the rise in intangible intensity to cheaper intangibles, with reductions in globalization frictions playing a secondary role. The large cross-effect from intangible costs onto multinational expansion reflects the interdependence between these margins: because intangible capital is nonrival, cheaper intangibles raise per-market profits and thereby the incentive to expand abroad. The model reproduces key cross-sectional and dynamic patterns not targeted in calibration, including the positive relationship between intangible intensity, multinational status, and productivity, and the within-firm increase in intangible intensity following multinational entry.

Our final contribution is to develop the implications of these forces for aggregate productivity measurement. We construct a bottom-up approach to measuring aggregate TFP that is consistent with the CES production technology and accounts for the nonrival nature of intangible capital. The approach aggregates plant-level productivity residuals using Domar weights, where each plant's residual is computed from the true production function rather than from a common-factor-share approximation. This distinction is consequential because nonrivalry and productivity heterogeneity jointly generate substantial variation in intangible intensity across plants: multinational firms deploy the same stock of intangibles across all production locations, and more productive firms optimally choose both greater intangible capital and broader multinational scope. Whether this heterogeneity matters for productivity measurement depends on the elasticity of substitution. When intangible and rival inputs are Cobb-Douglas, output elasticities are constant across plants and standard growth accounting is valid. When the elasticity exceeds one—as our calibrated value of 1.8 implies—differences in intangible intensity translate into differences in output elasticities, and imposing common factor shares generates a systematic measurement bias that grows with the dispersion of intangible intensity.

The quantitative implications for productivity measurement are striking. Our production-function-consistent measure shows that the intangible revolution and multinational expansion generate a near-zero change in aggregate TFP, with the small decline reflecting compositional

churning as marginal firms enter multinational status. The two standard approaches to measuring productivity err in opposite directions relative to this benchmark. The bottom-up Solow residual, which imposes common factor shares at the plant level, understates TFP by 4.5 percentage points, misattributing factor accumulation in a nonrival input as a decline in efficiency. The aggregate-residual methodology used by the BEA and BLS errs in the opposite direction, overstating TFP by 1.4 percentage points—recording positive productivity growth exceeding one percentage point in a model economy where the underlying technology is essentially unchanged. This upward bias reflects two channels of roughly equal magnitude: the use of common factor shares that fail to capture heterogeneous plant-level elasticities, and the booking of intangible capital in the headquarters country rather than where it is deployed in production. Because both channels are amplified by rising intangible intensity and expanding multinational production, the divergence between measured and true productivity widens as these trends continue.

Related Literature. Our paper is closely related to three strands of the literature. The first studies how intangible capital shapes firm size and market structure. Corrado, Hulten, and Sichel (2009) establish that intangible investment is as large as physical capital investment in the United States and that omitting it understates capital deepening and productivity growth. Crouzet, Eberly, Eisfeldt, and Papanikolaou (2022) show that the nonrivalry of intangibles explains why intangible-intensive firms are systematically larger and account for a rising share of aggregate activity; Jones and Tonetti (2020) formalize this logic, showing nonrivalry generates increasing returns and a wedge between private and social incentives. Aghion, Bergeaud, Boppart, Klenow, and Li (2023) and Klenow and Li (2025) connect these forces to rising concentration and declining dynamism. At the geographic margin, Kleinman (2025) shows that nonrivalry of headquarter services drives the spatial expansion of firms—a mechanism we extend to cross-border multinational expansion. We contribute by embedding these nonrivalry forces in a quantitative model of heterogeneous multinational firms where the global footprint is an endogenous choice, quantifying how changes in frictions to operating abroad translate into the joint distribution of firm scale, intangible intensity, and aggregate outcomes.

A second strand studies the determinants and aggregate implications of multinational production. Helpman, Melitz, and Yeaple (2004) establish that only the most productive firms select into multinational status—a proximity-concentration tradeoff robustly confirmed empirically and surveyed in Antràs and Yeaple (2014). Quantitative models show that gains from multinational production are large and comparable to gains from trade (Ramondo and Rodríguez-Clare, 2013), operating through cross-border diffusion of firm-specific technologies, with production networks further am-

plifying international technology transmission (Tintelnot, 2017). Alviarez, Cravino, and Ramondo (2023) measure this firm-specific transferable asset directly, isolating firm-embedded productivity—blueprints, management practices, and intangible capital—from country-embedded factors, and finding that the former accounts for roughly one-third of cross-country income differences. Our contribution is to endogenize the accumulation of this firm-specific, non-rival intangible asset, providing a unified structural account of the complementarity between intangible investment and multinational expansion.

A third literature concerns the implications of intangible capital and multinational production for the measurement of aggregate total factor productivity. Corrado et al. (2009) show that omitting intangible investment understates aggregate productivity growth. Crouzet and Eberly (2021) demonstrate that intangible mismeasurement and rising markups jointly amplify the downward bias in measured TFP growth. These measurement challenges are compounded in open economies by the cross-border nature of multinational production. McGrattan and Prescott (2009, 2010) show that the nonrival nature of technology capital creates a wedge between intangible capital booked at the headquarters country in national accounts and the capital actually deployed in production by foreign affiliates, generating systematic mismeasurement of aggregate TFP in open economies. Our paper contributes to this literature by developing a micro-founded framework that accounts for substitutability between inputs, reallocation, mismeasurement, and multinational production. In addition to formalizing aggregate TFP accounting through the model, a novel insight is that gross substitutability between intangible and rival inputs generates heterogeneous output elasticities across plants, so standard growth accounting—which imposes common factor shares—introduces a systematic bias amplified by reallocation toward intangible-intensive multinationals.

The remainder of the paper is organized as follows. Section 2 presents empirical evidence documenting the joint evolution of intangible intensity, multinational production, and profit shifting. Section 3 develops the theoretical framework and derives its main qualitative implications. Section 4 introduces the quantitative model and describes its equilibrium. Section 5 describes the calibration strategy and validates the model's key parameter estimates. Section 6 presents the main quantitative results. Section 7 studies the implications of the framework for aggregate productivity. Section 8 concludes.

2 Empirical Analysis

This section documents empirical evidence on the interconnection between intangible capital and multinational production. We proceed in four stages. First, we describe our data sources and

the construction of key variables. Second, we document secular trends in intangible intensity and multinational production over the period 1995–2023 and examine cross-sectional relationships between these phenomena using panel regressions. Third, we employ an event-study design to analyze the dynamic effects of multinational entry initiation on intangible intensity. Finally, we document the relationship between firm productivity and intangible intensity, providing evidence of selection into multinational status.

2.1 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.

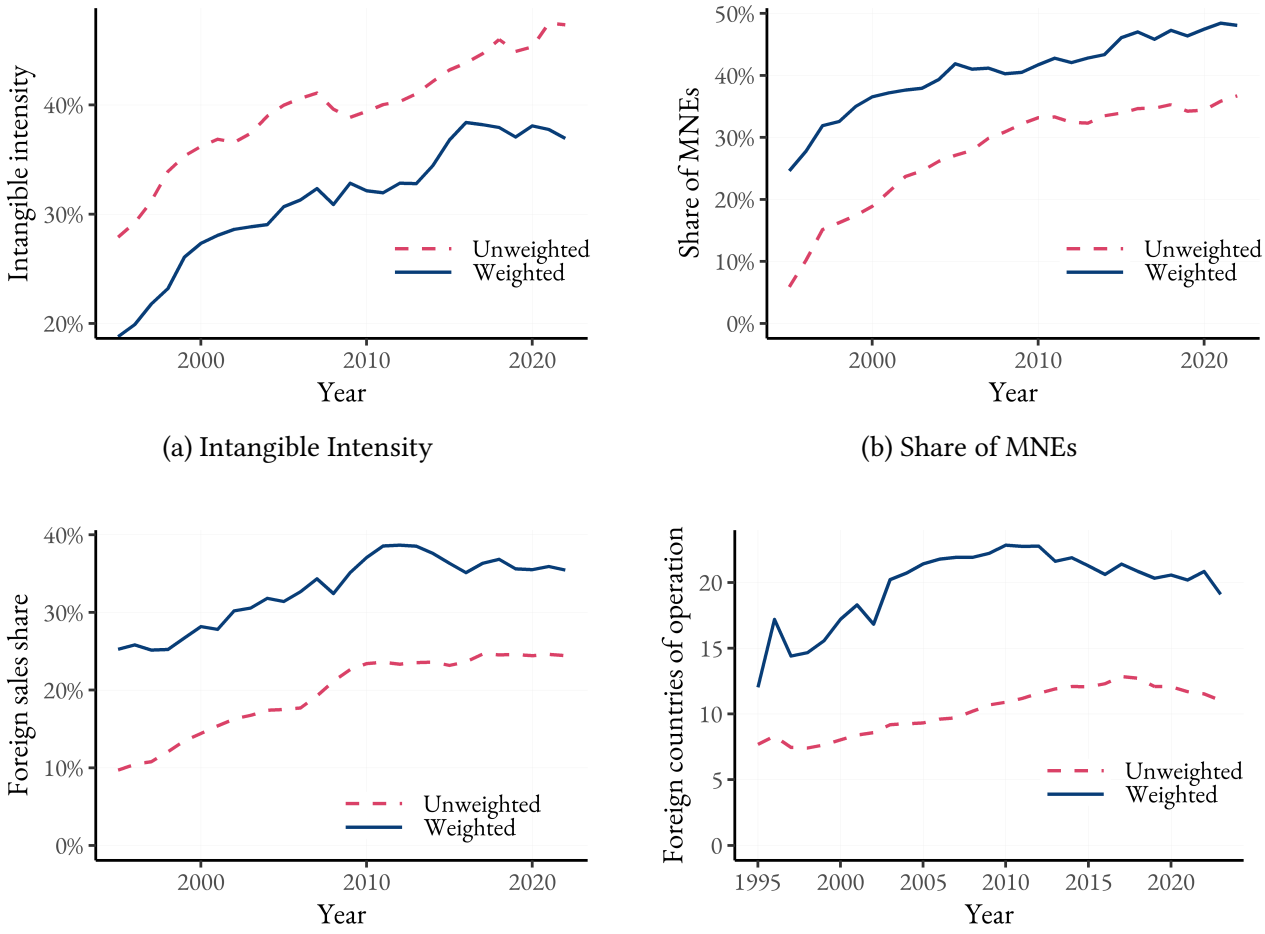
Measuring intangible capital. We construct intangible capital as the sum of balance-sheet intangibles and internally generated knowledge capital from capitalized R&D expenditures, following Peters and Taylor (2017)[1]. This measure captures the canonical nonrival inputs emphasized by our theory: patents, 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:

$$s_{jt}^{z} \equiv \frac{\text{Intangible Capital}_{jt}}{\text{Intangible Capital}_{jt} + \text{Tangible Capital}_{jt}}. \tag{1}$$

Measuring multinational activity. We link Compustat data 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. On the extensive margin, we classify a firm as a multinational enterprise (MNE) in year t if it reports at least one foreign subsidiary in its Exhibit 21 disclosure. On the intensive margin, we construct two mea-

[1]We elaborate on our construction of firm-level intangible capital in Appendix B.

Figure 1: Secular Trends of Intangible Intensity and MNE Activity: 1995 – 2023



(a) Intangible Intensity

(b) Share of MNEs

(c) Foreign sales share of MNEs

(d) Average Number of Subsidiary Countries

Notes: This figure displays the time series of key variables from 1995 to 2023 using Compustat North America data. Panel (a) shows intangible intensity, defined as the ratio of intangible assets to total assets. Intangible assets include balance-sheet intangible assets (externally acquired) and internally generated intangible capital constructed by capitalizing R&D expenditures using the perpetual inventory method following Peters and Taylor (2017). Total assets are defined as the sum of tangible (property, plant, and equipment) and intangible assets. Panel (b) shows the share of firms classified as multinational enterprises (MNEs), defined as firms reporting at least one foreign subsidiary in Exhibit 21 of SEC 10-K filings. Panel (c) displays the foreign sales share among MNEs, calculated as the ratio of foreign segment revenue to total consolidated revenue. Panel (d) displays the average number of foreign countries in which MNEs operate subsidiaries, as reported in Exhibit 21 of SEC 10-K filings. In all panels, dashed lines represent unweighted averages across firms, while solid lines represent revenue-weighted averages. The sample excludes firm-year observations with missing industry codes, or non-positive total revenue.

sures: (i) the number of distinct foreign countries in which the firm operates subsidiaries and (ii) the share of foreign segment revenue in total consolidated revenue.

2.2 Secular Trends

We document two secular trends over our sample period: the rise in intangible intensity and the expansion of multinational production. Figure 1 summarizes the trends using Compustat data.

Intangible intensity. Figure 1, Panel (a) displays the evolution of intangible intensity from 1995 to 2023. Average intangible intensity increased from 26% in 1995 to 47% in 2023, nearly a doubling over the sample period, with the revenue-weighted average exhibiting a similar trajectory (rising from 16% to 38%). This secular increase is consistent with extensive prior documentation in the literature (Corrado et al., 2009; Crouzet et al., 2022).

Multinational production. Figure 1, Panel (b) documents significant growth in multinational activity. The unweighted share of firms classified as MNEs increased from 7% in 1995 to approximately 30% by 2020 before declining slightly thereafter. The expansion is evident along the intensive margin as well: Panel (d) shows that the average number of foreign countries in which MNEs operate subsidiaries increased from approximately 7 in 1995 to around 10 by 2020, with revenue-weighted averages showing steeper growth. Panel (c) further demonstrates that the share of foreign sales in total MNE revenue has risen steadily, from 10% (unweighted) in 1995 to 22% by 2023. These patterns indicate a deepening of global integration beyond simply entering new markets.

2.3 More Global Firms Are More Intangible-Intensive

We examine the cross-sectional relationship between intangible intensity and multinational status using panel regressions that control progressively for year effects, firm revenue, and industry composition.

Table 1 reports results. Columns (1)–(3) examine the full sample using an indicator for MNE status. Column (1) includes only year fixed effects, yielding a coefficient of 0.09: MNEs have 9 percentage points higher intangible intensity than domestic firms. Adding a control for log revenue increases the coefficient to 0.12 (column 2), indicating that the unconditional correlation understates the conditional relationship: within revenue categories, MNEs are even more intangible-intensive.

Our preferred specification adds industry fixed effects (column 3). The baseline coefficient falls to 0.06 for the pre-2005 period, reflecting the fact that some of the raw MNE–domestic gap is explained by industry composition. The interaction with a post-2005 indicator is positive and statistically significant (0.05), implying that the intangible intensity gap between MNEs and domestic firms nearly doubled after 2005. This temporal pattern is consistent with the hypothesis that declining barriers to globalization and intangible investment have disproportionately benefited multinational firms.

Is the intangible gap purely about *being* an MNE, or does it scale with the extent of multinational

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 measures of multinational activity. The dependent variable in all columns is intangible intensity, defined as the ratio of intangible assets (including capitalized R&D) to total assets. "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 to MNEs only (to examine within-MNE variation in geographic scope). 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: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

activity? Columns (4)–(6) address this question by restricting the sample to MNEs and replacing the binary indicator with the log number of foreign countries in which the firm operates. The answer is clear: intangible intensity rises with geographic scope. In column (6), which includes both industry fixed effects and a revenue control, a 10% increase in the number of foreign countries is associated with about 0.3 percentage points higher intangible intensity before 2005 and 0.5 percentage points after. The median MNE in the sample operates in 8 countries, so expanding to one additional country (a roughly 12% increase) would imply about 0.4 percentage points before 2005 and 0.6 percentage points after. The fact that the gradient is present *within* MNEs rules out explanations based solely on selection into multinational status and points to a tight link between the scale of global operations and intangible capital use.

2.4 Dynamic Effects of MNE Entry

To examine the temporal relationship between multinational status and intangible intensity, we employ an event-study design that exploits variation in the timing of firms' transitions to MNE status.

We employ the interaction-weighted (IW) estimator, as introduced in Sun and Abraham (2021), to estimate the dynamic treatment effects of MNE status on intangible intensity for a balanced panel of firms spanning 1995 to 2023. Specifically, we focus on the impact of a firm first attaining MNE status. In the context of staggered treatment adoption, the IW estimator offers significant advantages. It mitigates contamination from treatment effects in other periods and accounts for treatment effect heterogeneity, an issue when applying the standard two-way fixed-effects (TWFE) event-study approach. To address variations in the duration that firms remain MNEs, we estimate treatment effects separately for cohorts of firms based on their initiation year of treatment. Specifically, we run the following regression

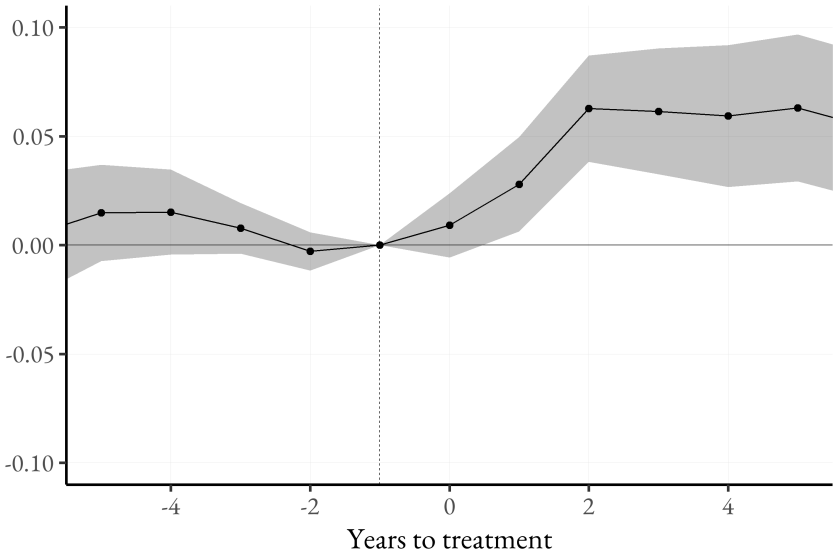
$$s_{jt}^{z} = \gamma_j + \delta_t + \sum_{e \notin C} \sum_{\tau \neq -1} \beta_{e,\tau} \cdot \mathbb{1}_{\{E_j = e\}} \cdot \mathbb{1}_{\{t - E_j = \tau\}} + \varepsilon_{jt}, \tag{2}$$

where E_j is the year when firm j becomes an MNE, and $\beta_{e,\tau}$ is the cohort-specific average treatment effect on the treated (CATT) τ periods from initial treatment for the cohort of firms first treated in year e. We use firms that later or never become MNEs as the control group C . We include all relative time indicators τ, excluding only the indicator for the period before treatment ($\tau = -1$) to address multicollinearity. The dynamic treatment effect estimate for each relative time period is given by the average of $\hat{\beta}_{e,\tau}$, weighted by the proportion of each cohort e in the relative time period τ. Treatment is absorbing: once a firm becomes an MNE, it remains treated in subsequent periods.

The IW estimator identifies CATT under two key assumptions. First, *parallel trends*: in the absence of treatment, treated and control firms would have experienced parallel evolution of intangible intensity. We assess this assumption by examining pre-treatment coefficients ($\tau < 0$); under parallel trends, these should be statistically indistinguishable from zero. Second, *no anticipation*: firms do not adjust their intangible intensity in anticipation of future treatment. This assumption may be violated if firms increase R&D investment prior to planned international expansion, in which case pre-treatment coefficients would be positive. We examine this possibility in our results.

A maintained assumption is that treatment timing is exogenous conditional on firm and year fixed effects. This is unlikely to hold exactly: firms endogenously choose when to expand internationally based on productivity shocks, market opportunities, and strategic considerations. We therefore interpret our estimates as descriptive of the *average experience* of firms before and after treatment, rather than as causal effects of treatment per se.

Figure 2: Event Study: Dynamic Effects of MNE Entry on Intangible Intensity



Notes: This figure displays dynamic treatment effect estimates from an event-study regression examining how intangible intensity evolves around firms' transitions to multinational status. The plotted coefficients are cohort-share-weighted averages of the $\hat{\beta}_{e,\tau}$ estimates from Equation (2). The omitted category is $\tau = -1$ (one year before treatment). Shaded bands represent 95% confidence intervals based on standard errors clustered at both the firm and year level. The flat pre-trends (coefficients near zero for $\tau < 0$) support the parallel trends assumption. The figure shows that MNE entry increases intangible intensity by approximately 6 percentage points within 2–4 years.

Figure 2 shows that becoming an MNE is associated with sustained increases in intangible intensity. MNE expansion increases intangible intensity by approximately 6 percentage points after 2-4 years, with the effect emerging around one year after treatment and stabilizing thereafter. The lagged responses in intangible intensity suggest the presence of adjustment costs in intangible investment.

The event-study results provide evidence of a treatment effect: conditional on a firm's prior trajectory, intangible intensity increases following multinational entry. This pattern is consistent with the nonrivalry mechanism emphasized by our model: once a firm incurs the fixed costs of foreign expansion, the marginal return to intangible investment rises because intangibles can be deployed across a larger production base.

However, the event-study estimates cannot distinguish this treatment effect from time-varying selection. If firms expand internationally precisely when they anticipate high returns to intangible investment (due to product opportunities, technological breakthroughs, or competitive pressures), the observed increases in intangible intensity may reflect these underlying shocks rather than a causal effect of MNE status per se.

Combined with the cross-sectional evidence from Section 2.3, the event studies suggest that both *selection* (i.e., high-productivity, high-intangible firms choose to become MNEs) and *treatment* (i.e., MNE entry raises intangible intensity) contribute to the observed correlation. Disentangling these channels requires the structural model developed in Section 3.

2.5 Productivity, Intangible Intensity, and Selection into Multinational Status

To further investigate the selection mechanism, we examine the relationship between firm-level productivity and intangible intensity. We estimate firm-level total factor productivity (TFP) using the Ackerberg, Caves, and Frazer (2015) estimator. Specifically, we assume a (log) Cobb-Douglas production technology:

$$y_{it} = \alpha + w_{it}\beta + k_{it}\gamma + \omega_{it} + \varepsilon_{it}$$

where y_{it} is log output (value-added); w_{it} is a vector of log free variables (employment headcount); k_{it} is a vector of state variables (tangible and intangible capital); ω_{it} is productivity; and ε_{it} is an idiosyncratic error term.[2] Following Levinsohn and Petrin (2003), we use materials (proxied by cost of goods sold) as the control variable to address the simultaneity of productivity and input choices.

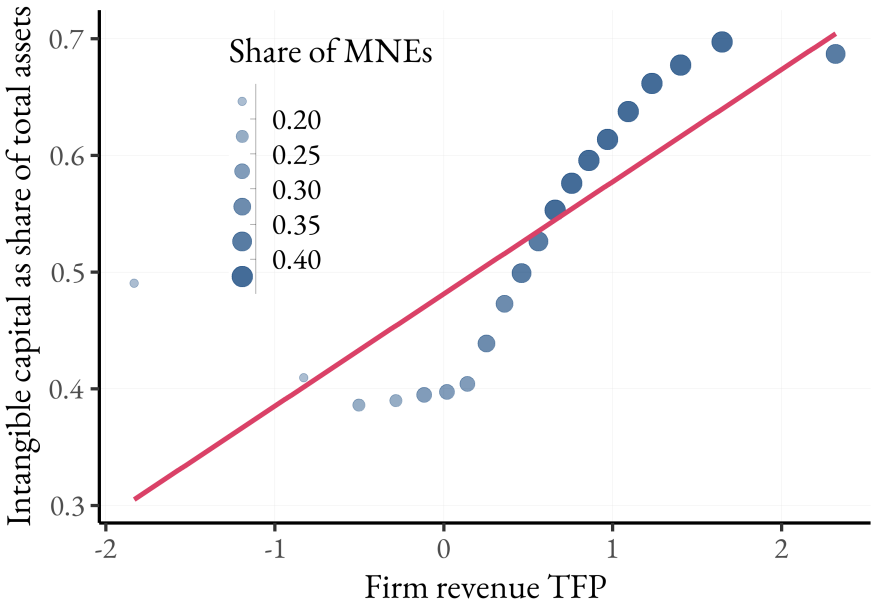
Figure 3 displays the relationship between firm TFP and intangible intensity, with bubble size proportional to the share of MNEs in each TFP vigintile. Appendix C provides corresponding regression evidence, confirming that these relationships are statistically robust to controls for industry composition and firm size. Three patterns emerge.

First, intangible intensity increases monotonically with productivity. Firms in the highest TFP vigintile exhibit intangible intensity approximately 20 percentage points above firms in the lowest vigintile. This positive gradient is consistent with models in which intangible capital is complementary to firm productivity, either because productive firms have higher marginal products of intangible investment, or because intangible investment itself raises measured productivity.

Second, MNE status increases sharply with productivity. The share of MNEs rises from ap-

[2]We use employment headcount rather than the wage bill as our labor input measure because Compustat's staff expense variable (XLR) has poor coverage, with missing values exceeding 70% of firm-year observations. This choice implicitly treats labor as homogeneous, so our revenue TFP estimates absorb cross-firm differences in workforce quality that would otherwise be captured by wage variation. If MNEs systematically employ higher-skilled workers than domestic firms, our estimates may overstate their productivity advantage. However, our focus on relative rankings rather than absolute magnitudes, combined with industry-year fixed effects that absorb systematic variation in workforce composition, mitigates this concern.

Figure 3: Selection into Multinational Status: Productivity and Intangible Intensity



Notes: This figure examines the relationship between firm-level productivity and intangible intensity, and how this relationship varies with multinational status. The x-axis displays firm revenue total factor productivity (TFP), estimated using the Ackerberg et al. (2015) control function approach. We assume a Cobb-Douglas production function in employment (free variable) and tangible plus intangible capital (state variables), using cost of goods sold as the proxy variable following Levinsohn and Petrin (2003). Firm-year TFP estimates are residualized by industry-year fixed effects and normalized to have mean zero. The y-axis displays intangible capital as a share of total assets. Each bubble represents a bin of firm-year observations grouped by TFP vigintiles. Bubble size is proportional to the share of MNEs within each bin, ranging from approximately 20% in the lowest TFP bins to over 40% in the highest bins. The fitted line (OLS) illustrates the strong positive relationship between productivity and intangible intensity.

proximately 20% in the lowest TFP bins to over 40% in the highest bins, consistent with the well-documented pattern that only the most productive firms find it profitable to bear the fixed costs of foreign expansion (Melitz, 2003; Helpman et al., 2004).

Third, the slope of the TFP-intangible intensity relationship appears steeper at higher productivity levels. This pattern is consistent with our theoretical result (Proposition 3) that cost reductions have larger effects for high-productivity firms: if intangible costs have fallen over time, the induced increase in intangible intensity should be concentrated among the most productive firms.

The productivity-based selection patterns documented in Figure 3 have important implications for interpreting the earlier cross-sectional and event-study results. The positive correlation between MNE status and intangible intensity (Table 1) reflects, at least in part, selection on productivity: firms that choose to become MNEs are more productive *and* more intangible-intensive, even absent any causal relationship between multinational status and intangible investment.

Our empirical analysis establishes four principal findings. First, intangible intensity and multinational production have both increased substantially since 1995, with particularly rapid growth

in the 2000s. Second, MNEs exhibit significantly higher intangible intensity than their domestic counterparts, with gaps that have widened over time. Third, firms transitioning to MNE status experience subsequent increases in intangible intensity. Fourth, high-productivity firms are disproportionately intangible-intensive and more likely to operate internationally, suggesting an important role for selection.

These patterns are consistent with the theoretical mechanisms we formalize in Section 3, but they do not permit sharp causal inference. The correlational nature of the evidence, combined with the clear presence of selection effects, motivates the structural approach we develop next.

3 Theory of Multinational Production and Intangible Investment

We develop a theory to study firms' decisions regarding multinational production and intangible investment. We use the framework to show how declining costs of multinational expansion and intangible investment affect firms' outcomes of interest, and how the effects of these cost reductions endogenously complement each other. At the heart of the theory is the interplay between the non-rivalry of intangible capital and multinational expansion: when intangible investment costs fall, firms have stronger incentives to scale globally; when expansion frictions fall, the returns to investing in intangibles rise. The model is deliberately parsimonious, delivering transparent comparative statics and a clean mapping to empirical moments.

3.1 Environment

Suppose a world that consists of a set of identical productive countries, $\mathbb{J} = [0, \bar{J}]$. In each country there exists a set of firms that differ in productivity $a \in \mathbb{A}$. 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 serve foreign markets, firms set up subsidiaries abroad, subject to a multinational expansion cost that depends on the number of countries served, $C^J(J)$. Firms also invest in intangible capital z at marginal cost MC. Intangible capital is non-rival and can be used in all subsidiaries. Each subsidiary produces a final good using the local rival input l and the firm-level non-rival input z. All firms take the input prices of the rival factor w, the marginal cost of intangible capital MC, and the output price p as given.

The profit-maximization problem of a firm with productivity a is

$$\pi(a) = \max_{l,z,J} \left\{ \pi^G(a,l,z,J) - MC \cdot z - C^J(J) \right\}$$

$$\text{where} \quad \pi^G(a,l,z,J) = J \times \pi(a,l,z)$$

$$\pi(a,l,z) = p \times aF(z,l) - wl.$$

We define intangible intensity as the expenditure share on intangible capital, $s^z \equiv MC \cdot z/(Jwl + MC \cdot z)$, which is the model counterpart of the empirical measure constructed in Section 2.

We make the following assumptions on the production technology and the expansion cost.

Assumption 1. *We make the following parametric form assumptions on the production and cost functions:*

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

$$F(z,l) = \left[\mu z^{\frac{\sigma-1}{\sigma}} + (1-\mu) l^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma\alpha}{\sigma-1}}$$

   where α is the returns-to-scale parameter, σ is the elasticity of substitution between z and l, and μ is the input weight of z in the CES production function. We assume $\alpha < 1$, $\sigma > 1$, and the weight on the rival input, $1-\mu$, is sufficiently large.

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

$$C^J(J) = \bar{C}^J J^\gamma$$

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

Assuming DRS technology ($\alpha < 1$) is standard in an environment of perfect competition and heterogeneous firms. Assuming $\sigma > 1$ implies that the firm substitutes toward the relatively cheaper production input. This gross substitutability assumption is supported by the empirical literature that estimates the elasticity of substitution between intangible capital and rival factors such as labor and tangible capital (Jerbashian, 2024a). The assumption on μ requires the rival

factor to account for a non-trivial input share, and it is needed for the complementarity results we establish below. We assume that the cost of expansion is convex in the number of foreign subsidiaries, which reflects information frictions and span-of-control costs of managing a greater number of establishments (Oberfield, Rossi-Hansberg, Sarte, and Trachter, 2024). The assumption on γ ensures that no firm expands indefinitely.

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 β is defined as $\varepsilon_{y,\beta} = \frac{\partial \ln y}{\partial \ln \beta}$.*
2. *The combined elasticity of an endogenous variable y with respect to variables β_1 and β_2 is defined as*

$$\varepsilon_{y,(\beta_1,\beta_2)} = \varepsilon_{y,\beta_1} + \varepsilon_{y,\beta_2} + \frac{\partial^2 \ln y}{\partial \ln \beta_1 \partial \ln \beta_2}.$$

With the functional form assumptions and the definitions above, we can derive the endogenous variables of interest from the first-order conditions of the firm as follows.

$$s^z = \left[\left(\frac{1-\mu}{\mu} \right)^{\sigma} \left(\frac{w}{MC} \right)^{1-\sigma} J^{1-\sigma} + 1 \right]^{-1} \tag{3}$$

$$J = \left(\frac{1}{\gamma \bar{C}^J} \pi \left(a, l, z \right) \right)^{\frac{1}{\gamma - 1}}. \tag{4}$$

These expressions highlight the two-way interaction between multinational expansion and intangible investment. First, multinational expansion incentivizes intangible capital deepening at the firm level when $\sigma > 1$. Second, more intangible investment enhances per-market profits π, encouraging a greater extent of expansion.

3.2 Theoretical Results

Comparative Statics on Cost Reduction

We first establish a set of comparative statics on how reductions in intangible cost MC and expansion cost $\bar{C}^J$ affect intangible intensity s^z and multinational expansion J.

Proposition 1. *When intangible cost MC or expansion cost $\bar{C}^J$ falls, the following statements hold under the assumptions:*

1. *all firms increase in intangible intensity, i.e.* $\varepsilon_{s^z,\beta} < 0$, $\beta \in \{MC, \bar{C}^J\}$;

2. *all firms expand into more foreign countries, i.e.* $\varepsilon_{J,\beta} < 0$, $\beta \in \{MC, \bar{C}^J\}$.

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

Lemma 2. *We have* $\frac{\partial s^z}{\partial \beta} = \frac{\partial s^z}{\partial \beta}|_J + \frac{\partial s^z}{\partial J}\frac{\partial J}{\partial \beta}$, $\beta \in \{MC, \bar{C}^J\}$. *Endogenous multinational expansion amplifies the effect of cost reductions on the intangible intensity of a firm, i.e.* $\frac{\partial s^z}{\partial J}\frac{\partial J}{\partial \beta} < 0$.

The effect of expansion J on intangible intensity s^z can be seen directly in equation (3). Intuitively, cost reductions affect intangible intensity partly through their effect on expansion J. As intangible capital is non-rival, an increase in J raises the firm-level return to z. When the production inputs are gross substitutes, $\sigma > 1$, firms optimally substitute toward the non-rival intangible as its marginal return increases. This investment effect, in turn, raises profits in each market and further fuels expansion.

Complementarity

Next, we formalize the complementary effects of reductions in MC and $\bar{C}^J$ on s^z and J.

Proposition 2. *For* $y = \{s^z, J\}$, *reductions in* $\bar{C}^J$ *and* MC *are complementary: their joint effect on* y *is super-additive,*

$$|\varepsilon_{y,(\bar{C}^J,MC)}| > |\varepsilon_{y,\bar{C}^J}| + |\varepsilon_{y,MC}|.$$

The inequality is expressed in absolute values because the elasticities are negative: reductions in the costs increase s^z and J. The proposition states the complementary effect of expansion cost $\bar{C}^J$ and intangible investment cost MC. Such complementarity stems from the interdependence of multinational expansion and intangible investment. For the expansion margin, the effect of lowering expansion costs operates through the induced change in per-market profits π. When intangible investment is cheaper, firms raise their intangible capital to a greater extent as they enter additional markets, which boosts π and amplifies the expansion response. For the intangible-intensity margin, the effect of lowering investment cost operates through the induced expansion in J. When expansion is cheaper, firms open more foreign subsidiaries and this further boosts the return to intangible investment.

Firm Heterogeneity

The last set of theoretical results concerns the role of firm heterogeneity. We begin by showing comparative statics of firm productivity a on s^z and J.

Proposition 3. *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.* $\varepsilon_{J,a} > 0$.

The proposition states that more productive firms choose higher intangible intensity and a larger global footprint. This selection effect comes from the monotonicity of firm profits in productivity: holding fixed production inputs, more productive firms earn greater profits in a market. As a result, high-productivity firms are able to expand into more markets in the face of convex expansion costs, which raises their intangible intensity as shown in Proposition 1. These theoretical predictions align with our empirical findings: in Appendix C, we document that higher-productivity firms exhibit greater intangible intensity (Table 8) and broader multinational scope (Table 7), with these relationships robust to controlling for industry and firm size.

Finally, we show that the effects of cost reductions are greater for high-productivity firms.

Proposition 4. *The following statements are true under the assumptions for* $y = \{s^z, J\}$*:*

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

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

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

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

The proposition implies that cost reductions have larger effects for high-productivity firms. The intuition is the same as above: the impact of a cost decline hinges on the firm's capacity to expand, which rises with productivity. Higher productivity effectively lowers firm-specific expansion costs, or equivalently raises per-market returns, so firm-level efficiency complements economy-wide cost reductions, giving rise to stronger global expansion and greater intangible intensity among high-productivity firms.

4 Quantitative Theory

This section develops a dynamic, multi-country general equilibrium model designed to quantify the joint evolution of intangible investment, and multinational expansion. The model embeds the core mechanisms highlighted in the analytical theory—nonrival intangible capital and endogenous multinational scope—into a tractable quantitative framework with heterogeneous firms and forward-looking investment decisions. The objective is twofold: first, to discipline the strength of the key complementarities using firm-level data, and second, to conduct counterfactual experiments that decompose the observed trends into technological change and globalization.

4.1 Households.

Each region i has a representative household with preferences over sequences of consumption, $\{C_{it}\}_{t=0}^{\infty}$, and labor supply, $\{L_{it}\}_{t=0}^{\infty}$, given by

$$\sum_{t=0}^{\infty} \beta^t \left[\log\left(\frac{C_{it}}{N_i}\right) + \psi_i \log\left(1 - \frac{L_{it}}{N_i}\right) \right]. \tag{5}$$

Households choose consumption, labor supply, tangible investment, $\{X_{it}\}_{t=0}^{\infty}$, and internationally-traded bonds, $\{B_{it+1}\}_{t=0}^{\infty}$ to maximize utility subject to a sequence of budget constraints,

$$P_{it}[C_{it} + K_{it+1} - (1-\delta)K_{it}] + P_{bt}B_{it+1} = W_{it}L_{it} + R_{it}K_{it} + B_{it} + D_{it}, \tag{6}$$

taking the wage, W_{it}, the rental rate, R_{it}, the bond price, P_{bt}, dividends, D_{it}, and lump-sum transfers T_{it} as given.

4.2 Firms

Firms are the central decision-making units in the model. They are heterogeneous in productivity and dynamically choose intangible investment and multinational expansion in response to prices and policy frictions. Each firm is headquartered in a single country but may operate production units in multiple foreign locations, allowing intangible capital developed at headquarters to be deployed globally. We describe the firm problem in stages, beginning with production technology and intangible accumulation, followed by multinational expansion, and the firm's static and dynamic optimization problems.

Production technology. Firms are headquartered in a home country i and may operate production units in a set of foreign countries j. A firm with productivity a and intangible capital stock z produces output in country j according to

$$f_{ij}(a, z, k_j, \ell_j) = a \cdot \mathbf{1}\{i \neq j\}(1-\eta)\left[\mu z^{\frac{\sigma-1}{\sigma}} + (1-\mu)\left(k_j^{\alpha}\ell_j^{1-\alpha}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\theta}, \tag{7}$$

where k_j and ℓ_j denote tangible capital and labor hired locally in country j. Intangible capital z is a nonrival input that can be deployed simultaneously across all production locations of the firm, while tangible capital and labor are rival and location-specific. The parameter α governs the labor share within the rival input bundle, μ controls the importance of intangibles in production, and $\sigma > 1$ is the elasticity of substitution between intangible capital and the composite rival input. The parameter $\eta \in (0, 1)$ captures variable barriers to foreign production, which reduce effective productivity in foreign affiliates relative to the headquarters location. Finally, $\theta < 1$ governs returns to scale at the plant level, ensuring decreasing returns to production at each location.

Intangible capital accumulation. Intangible capital is accumulated at the firm level and is *nonrival across production locations.* At the beginning of each period, a firm enters with an existing stock of intangible capital z and chooses new intangible investment x, denominated in units of the home country's final good. Intangible capital evolves according to

$$z' = (1-\delta)z + x,$$

where $\delta \in (0, 1)$ is the depreciation rate of intangible capital. Investment in intangibles is subject to a constant marginal cost MC and a convex adjustment cost, so that total investment expenditures equal

$$\left[MC \cdot x + \frac{\phi}{2}x^2\right],$$

where $\phi > 0$ governs the degree of adjustment frictions. Adjustment costs smooth the dynamics of intangible accumulation and generate gradual responses of intangible intensity to shocks, consistent with the delayed investment responses observed in the data. Because intangible capital is nonrival, a higher stock of z raises productivity simultaneously across all domestic and foreign production units, implying that the return to intangible investment is increasing in the firm's multinational scope.

Multinational expansion. Firms may operate production units in multiple foreign countries in addition to their headquarters location. Let $n \in \{0, 1, \ldots, \bar{J} - 1\}$ denote the number of foreign

affiliates a firm operates at the beginning of a period. After observing current productivity and intangible capital, the firm draws an idiosyncratic fixed cost of expanding to an additional foreign country. The cost of expanding from n to $n+1$ affiliates is given by

$$\bar{C}_J = \bar{C}_J^0 - \varepsilon, \qquad \varepsilon \sim \text{Gumbel}(0, \sigma_{C_J}),$$

where $\bar{C}_J^0$ governs the average level of expansion costs and σ_{C_J} controls their dispersion. After observing the draw, the firm chooses whether to remain at n affiliates or expand to $n+1$ by paying the fixed cost. Expansion is irreversible within the period and subject to the upper bound $\bar{J}$ on the number of foreign locations.

Multinational expansion increases the scale over which the firm can deploy its intangible capital. Because intangible capital is nonrival, a higher number of affiliates raises the return to intangible investment by allowing the same stock of z to be used simultaneously across more production units. This mechanism generates endogenous selection into multinational status: firms with higher productivity or higher intangible capital are more likely to find expansion profitable. The stochastic cost structure ensures smooth aggregate responses of multinational activity to changes in technology, globalization frictions, and policy parameters.

Static profit maximization. Given productivity a, intangible capital z, and the number of foreign affiliates n, the firm chooses local production inputs to maximize global profits. Let country i denote the headquarters location and $j = 1, \ldots, n$ index foreign affiliates. In each location, the firm rents tangible capital and hires labor at competitive prices. Static global profits are given by

$$\begin{aligned} \pi(a, z, n) \quad = \quad & \max_{\{k_j, n_j\}_{j=1}^n, k_i, n_i} \Bigg[P_i f_{ii}(a, z, k_i, \ell_i) - W_i \ell_i - (R_i + \delta) k_i + \\ & + \quad \sum_{j=1}^{n} \left(P_j f_{ij}(a, z, k_j, \ell_j) - W_j \ell_j - (R_j + \delta) k_j \right) \Bigg] \end{aligned} \tag{8}$$

where P_j and W_j denote the price of the final good and the wage in country j, and δ is the depreciation rate of tangible capital.

Within the period, the firm takes (a, z, n) as given and chooses $\{k_j, \ell_j\}_{j=0}^n$ and λ to maximize profits. Optimal input choices equalize marginal revenue products to factor prices in each location. Static profits $\pi(a, z, n)$ summarize the payoff-relevant state for the firm and enter the dynamic problem governing intangible investment and multinational expansion decisions.

Dynamic program. Let the firm's state at the beginning of a period be (a, z, n), where a is productivity, z is the stock of nonrival intangible capital, and $n \in \{0, 1, \ldots, \bar{J} - 1\}$ is the number of foreign affiliates. Given the within-period static problem, the firm earns optimized profits $\pi(a, z, n)$. The firm then chooses intangible investment x (and, equivalently, next period's stock $z' = (1 - \delta)z + x$), and faces an expansion opportunity governed by the stochastic fixed cost described above. Let $V(a, z, n)$ denote the value function and define the post-investment continuation value

$$W(a, z', n) \equiv \rho_a V(a, z', n) + (1 - \rho_a) \int V(a', z', n)\, dG(a'), \tag{9}$$

where the firm survives to the next period with probability ψ and, conditional on survival, keeps its current productivity with probability ρ_a and otherwise draws a' from distribution G.

Under the i.i.d. Gumbel structure for the expansion cost draw, the discrete choice between remaining at n affiliates and expanding to $n + 1$ affiliates admits a closed-form expression for the expected continuation value. Let

$$\tilde{W}(a, z', n) = \sigma_{C_J} \log\left[\exp\left(\frac{W(a, z', n)}{\sigma_{C_J}}\right) + \mathbf{1}\{n < \bar{J} - 1\} \exp\left(\frac{W(a, z', n + 1) - \bar{C}_J}{\sigma_{C_J}}\right)\right], \tag{10}$$

which represents the expected value, conditional on (a, z'), of optimally choosing whether to expand the firm's multinational footprint in the current period.

Using this representation, the firm's value function can be written compactly as

$$V(a, z, n) = \max_x \left\{\pi(a, z, n) - \left[MC \cdot x + \tfrac{\phi}{2} x^2\right] + \beta\psi\, \tilde{W}(a, z', n)\right\}. \tag{11}$$

Intuitively, the firm chooses intangible investment to maximize current payoffs plus the expected continuation value, where the option value of expanding from n to $n + 1$ affiliates is smoothed by the extreme-value dispersion parameter σ_{C_J}. The resulting probability of expansion is

$$p(a, z', n) = \psi\, \mathbf{1}\{n < \bar{J} - 1\}\, \frac{\exp\left(\frac{W(a, z', n+1) - \bar{C}_J}{\sigma_{C_J}}\right)}{\exp\left(\frac{W(a, z', n)}{\sigma_{C_J}}\right) + \mathbf{1}\{n < \bar{J} - 1\} \exp\left(\frac{W(a, z', n+1) - \bar{C}_J}{\sigma_{C_J}}\right)}, \tag{12}$$

which is increasing in the net gain from operating one additional affiliate. This dynamic structure delivers gradual intangible adjustment through ϕ and smooth extensive-margin responses through

the Gumbel-logit expansion technology.

4.3 Aggregation and Equilibrium

We close the model by aggregating firm-level decisions and defining a stationary recursive equilibrium. Aggregate outcomes are obtained by integrating firms' optimal policies over the stationary distribution of firm states, accounting for multinational production and cross-border factor usage. We then characterize the law of motion for the firm distribution and define equilibrium conditions governing prices, allocations, and market clearing across countries.

Aggregation. Let $\Psi_i(a, z, n)$ denote the stationary joint distribution of firms headquartered in country i over productivity a, intangible capital z, and the number of foreign affiliates n. Given firms' optimal static and dynamic policies, aggregate quantities in each country are obtained by integrating firm-level decisions over this distribution.

Labor market clearing in country i requires that total labor supplied equals labor demand from both domestic firms headquartered in i and foreign affiliates operating in i:

$$\bar{L}_i = \int \ell_{ii}(a, z, n)\, d\Psi_i(a, z, n) + \sum_{j \neq i} \int \ell_{ji}(a, z, n)\, d\Psi_j(a, z, n),$$

where $\ell_{ji}(a, z, n)$ denotes labor hired in country i by a firm headquartered in country j.

Similarly, capital market clearing requires

$$K_i = \int k_{ii}(a, z, n)\, d\Psi_i(a, z, n) + \sum_{j \neq i} \int k_{ji}(a, z, n)\, d\Psi_j(a, z, n),$$

where $k_{ji}(a, z, n)$ denotes tangible capital rented in country i by firms headquartered abroad.

Aggregate output in country i is given by

$$Y_i = \int f_{ii}(a, z, k_i, \ell_i)\, d\Psi_i(a, z, n) + \sum_{j \neq i} \int f_{ji}(a, z, k_i, \ell_i)\, d\Psi_j(a, z, n).$$

The final good in each country is used for consumption, tangible investment, intangible investment, adjustment costs, and multinational expansion costs. Resource feasibility implies

$$Y_i = C_i + K_i' - (1 - \delta)K_i + \int \left[\frac{x_i(a, z, n)}{MC} + \frac{\phi}{2} x_i(a, z, n)^2 + p_i(a, z, n)\bar{C}_J \right] d\Psi_i(a, z, n),$$

where $x_i(a, z, n)$ denotes intangible investment and $p_i(a, z, n)$ the probability of multinational expansion.

Laws of motion. The evolution of the firm distribution is governed by firms' optimal policies for intangible investment and multinational expansion, together with exogenous productivity shocks and firm exit. Let (a, z, n) denote the firm's state at the beginning of a period. After choosing intangible investment x, the firm's intangible capital evolves deterministically according to

$$z' = (1 - \delta)z + x(a, z, n).$$

Productivity follows a Markov process. Conditional on survival, a firm retains its current productivity with probability ρ_a and draws a new productivity level a' from distribution $G(a')$ with probability $1 - \rho_a$. Firms exit exogenously at the end of the period with probability $1 - \psi$.

Multinational expansion occurs after intangible investment and before the next period's state is realized. A firm operating n foreign affiliates expands to $n + 1$ affiliates with probability $p(a, z', n)$ implied by the logit expansion rule, provided $n < \bar{J} - 1$. Otherwise, the firm remains at n. Expansion is irreversible within the period and subject to the upper bound $\bar{J}$.

Let $\Psi_i(a, z, n)$ denote the distribution of firms headquartered in country i. The law of motion for this distribution is given by

$$\Psi_i'(A, Z, n) = \psi \int \left[(1 - p(a, z', n))\, \mathbb{P}(a' \in A)\, \mathbf{1}\{z' \in Z\}\right] d\Psi_i(a, z, n),$$

for $n = 0, \ldots, \bar{J} - 1$, and

$$\Psi_i'(A, Z, n + 1) = \psi \int \left[p(a, z', n)\, \mathbb{P}(a' \in A)\, \mathbf{1}\{z' \in Z\}\right] d\Psi_i(a, z, n),$$

for $n = 0, \ldots, \bar{J} - 2$, where $\mathbb{P}(a' \in A) = \rho_a \mathbf{1}\{a \in A\} + (1 - \rho_a)G(A)$. Exiting firms are replaced by entrants with $z = 0$, $n = 0$, and productivity drawn from $G(a)$, ensuring a stationary mass of firms.

Equilibrium. A stationary recursive equilibrium consists of: (i) firm policy functions for intangible investment $x(a, z, n)$, and multinational expansion probabilities $p(a, z', n)$; (ii) value functions $V(a, z, n)$ and continuation values $W(a, z, n)$; (iii) stationary firm distributions $\{\Psi_i(a, z, n)\}_i$; and (iv) prices $\{W_i, R_i, P_i\}_i$ and aggregate quantities $\{C_i, K_i, Y_i\}_i$, such that:

1. Given prices, taxes, and policies, firm policy functions solve the static profit maximization problem and the dynamic program.

2. Firm distributions $\{\Psi_i\}$ are stationary and satisfy the laws of motion induced by firm policies, productivity shocks, entry, and exit.

3. Labor and capital markets clear in every country i.

4. Goods markets clear and resource constraints hold in every country.

The equilibrium is symmetric across countries and time-invariant.

5 Calibration and Validation

We calibrate the stationary equilibrium of the model to match key features of U.S. publicly listed firms over the period 2005–2022, using data from Compustat and SEC subsidiary disclosures described in Section 2. The model period is one year. Our calibration proceeds in two stages. First, we set a subset of parameters externally based on standard values in the literature or direct data counterparts. Second, we jointly calibrate the remaining parameters to match a set of targeted moments that discipline the model's extensive and intensive margins of multinational activity, intangible investment, and firm dynamics.

Externally set parameters. Panel A of Table 2 reports the externally calibrated parameters. We set the returns-to-scale parameter at the plant level to $\theta = 0.70$, consistent with estimates of decreasing returns in heterogeneous-firm models. The capital share within the rival input composite is set to $\alpha_k = 0.05$, implying an overall labor share of $\theta(1-\alpha_k) = 0.665$, in line with the aggregate labor share in the U.S. economy. The depreciation rate for tangible capital is set to $\delta_k = 0.10$, a standard annual value.

For the productivity process, we assume that firm-level productivity is drawn from a Pareto distribution with tail parameter $\kappa_a = 4.0$, a value commonly used to match the observed right tail of the firm size distribution (Helpman et al., 2004; Chaney, 2008). Conditional on survival, a firm retains its current productivity draw with probability $\rho_a = 0.70$ and otherwise draws a new productivity level from the stationary distribution. The annual firm survival probability is set to $\psi = 0.90$, consistent with observed exit rates among U.S. public firms. Finally, we set the number of possible production locations to $J = 2$ (one domestic and one foreign), capturing the binary extensive margin of multinational status. This simplification keeps the state space tractable while

Table 2: Model Parameters

Parameter	Description	Value
Panel A: Set Externally		
θ	Returns to scale	0.700
α	Capital share in rival composite	0.050
δ_k	Depreciation rate (tangible capital)	0.100
κ_a	Pareto tail parameter (productivity)	4.000
ρ_a	Productivity persistence	0.700
ψ	Survival probability	0.900
J	Number of locations (domestic + foreign)	2
Panel B: Calibrated		
σ_ε	Expansion cost shock dispersion (Gumbel)	0.003
ϕ_2	Intangible adjustment cost (quadratic)	0.000
σ	Elasticity of substitution (intangibles vs. rival)	1.843
μ	CES weight on intangible capital	0.059
η	Variable FDI barrier	0.123
$\bar{C}_J$	Mean fixed expansion cost	0.087
$z_{\min}$	Intangible capital lower bound	0.0010

Note: Implied overall labor share: $\theta(1-\alpha_k) = 0.665$.

preserving the key economic forces: the choice of whether to operate abroad and the interaction between that choice and intangible investment.

Jointly calibrated parameters. Panel B of Table 2 reports the seven parameters calibrated internally: the elasticity of substitution between intangible and rival inputs (σ), the CES weight on intangible capital (α_z), the variable FDI barrier (η), the mean fixed cost of multinational expansion ($\bar{C}_J$), the dispersion of expansion cost shocks (σ_ε), the quadratic intangible adjustment cost (φ_2), and the lower bound on intangible capital ($z_{\min}$). These parameters are chosen jointly to minimize the distance between model-generated moments and their data counterparts, where the targeted moments span both the extensive and intensive margins of firm behavior.

Extensive margin moments. We target three moments that discipline the scale and dynamics of multinational activity. The MNE transition rate—the fraction of domestic firms that become multinationals in a given year—averages 5.1% in the data, providing discipline on the flow into multinational status. The MNE share—the fraction of firms operating at least one foreign subsidiary—averages 50.2% over the 2005–2022 period. The ratio of foreign to domestic sales among MNEs averages 0.42, capturing the intensive margin of multinational production conditional on entry.

Intensive margin moments. We target three moments that discipline intangible investment behavior. Average intangible intensity across all firms is 33.8%, measured as the ratio of intangible

Table 3: Targeted Moments: Data and Model

Moment	Data	Model
Extensive Margin		
MNE transition rate	0.051	0.066
MNE share	0.502	0.460
Foreign/domestic sales ratio	0.420	0.425
Intensive Margin		
Average intangible intensity	0.338	0.373
MNE–nonMNE intensity gap	0.146	0.101
Median within-firm s.d. Δ investment rate	0.204	0.283
Cross-Sectional Discipline		
$\log(\text{sales})$–intensity slope	1.578	1.557
Dynamic Response to Becoming MNE		
Event study effect $(t+1)$	0.048	0.056
Event study effect $(t+2)$	0.062	0.072
Event study effect $(t+3)$	0.050	0.078
Event study effect $(t+4)$	0.056	0.081
Event study effect $(t+5)$	0.030	0.082

capital (including capitalized R&D) to total assets. The gap in intangible intensity between MNEs and non-MNEs averages 14.6 percentage points, a key moment for identifying the elasticity of substitution σ, as discussed below. The median within-firm standard deviation of changes in the intangible investment rate is 0.204, which disciplines the magnitude of adjustment costs in intangible accumulation.

Cross-sectional discipline. We also target the slope coefficient from a regression of intangible intensity on log sales, which equals 1.578 in the data. This moment captures the cross-sectional gradient between firm scale and intangible intensity and provides additional discipline on how the model maps productivity heterogeneity into variation in intangible investment across firms.

Dynamic response. Finally, we target the event-study coefficients documenting the evolution of intangible intensity following the transition to MNE status. The estimated treatment effects at horizons $t+1$ through $t+5$ after multinational entry range from 3.0 to 6.2 percentage points in the data. These moments discipline the model's dynamic adjustment path and help pin down the speed at which firms increase intangible investment after expanding abroad.

Model fit. The model provides a reasonable fit to the targeted moments (Table 3). On the extensive margin, the model generates an MNE share of 46.0% (versus 50.2% in the data), an MNE

transition rate of 6.6% (versus 5.1%), and a foreign-to-domestic sales ratio of 0.425 (versus 0.420). On the intensive margin, the model produces average intangible intensity of 37.3% (versus 33.8%) and an MNE–non-MNE intensity gap of 10.1 percentage points (versus 14.6 percentage points). The within-firm volatility of intangible investment is somewhat higher in the model (0.283) than in the data (0.204), reflecting the difficulty of simultaneously matching the level and dynamics of intangible adjustment. The cross-sectional slope of intangible intensity on log sales is 1.557 in the model, closely matching the data value of 1.578.

The event-study coefficients are broadly consistent with the data, though the model generates a somewhat more persistent response than the empirical estimates. The model produces effects of 5.6 percentage points at $t+1$, rising to 8.1 percentage points by $t+4$, compared to data values of 4.8 and 5.6 percentage points at the same horizons. The model's slightly larger and more persistent response reflects the absence of firm-level shocks that, in the data, may attenuate the measured treatment effect over longer horizons.

Identification. We assess local identification of the calibrated parameters using the sensitivity diagnostics of Andrews, Gentzkow, and Shapiro (2017). The weighted Jacobian matrix $G'WG$ has full numerical rank (7 out of 7), confirming that all seven calibrated parameters are locally identified.

The key parameter for the model's mechanisms is the elasticity of substitution σ, which is primarily identified by the MNE–non-MNE intangible intensity gap and the cross-sectional slope of intangible intensity on log sales. Because MNEs are selected on productivity, the gap between their intangible intensity and that of domestic firms reflects how strongly firms substitute toward intangible capital as their scale of operations expands. Figure 4 illustrates that σ traces out a smooth, monotone mapping between the elasticity and the intensity gap, ensuring that the moment condition has a unique solution in the empirically relevant range. The calibrated value of $\sigma = 1.843$ is closely aligned with the IV estimate of $\hat{\sigma} = 1.81$ reported in Section 5.1, providing independent external validation based on exogenous variation in FDI barriers.

5.1 External Validation of the Elasticity of Substitution σ

The calibration in the previous section pins down σ, jointly with other model parameters, by matching the intangible intensity gap between MNEs and domestic firms. To validate our σ estimate, this section exploits the structural relationship between intangible intensity and multinational reach alongside an instrumental variable approach. The structural equation (3) implies the

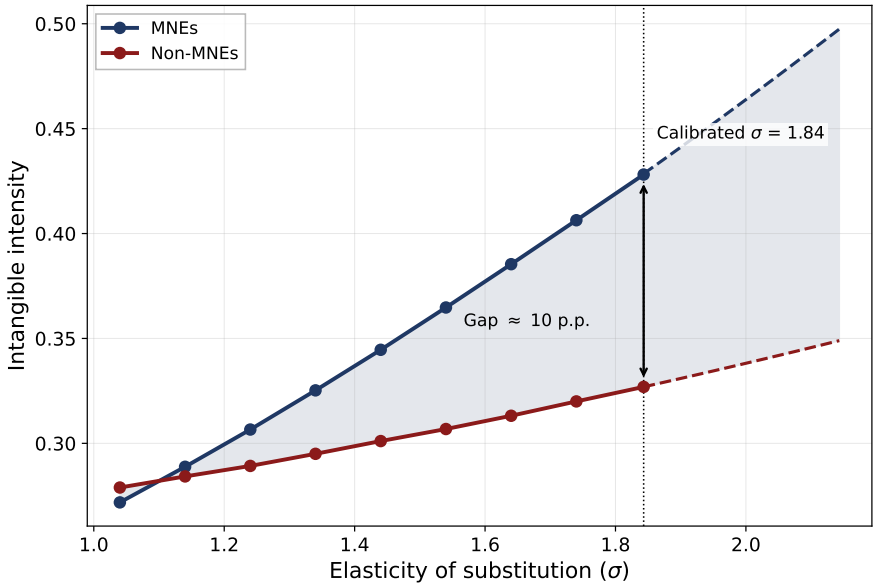


Figure 4: Intangible intensity as a function of σ

following empirical relationship between s_{it}^z and J_{it}:

$$\log\left(\frac{1-s_{it}^z}{s_{it}^z}\right) = (1-\sigma)\log J_{it} + \alpha_i + \delta_t + \varepsilon_{it}, \tag{13}$$

where i indexes a firm and t indexes year, s_{it}^z is the intangible capital share, J_{it} is the number of countries in which the firm operates, α_i is the firm fixed effect that absorbs fixed firm characteristics, δ_t is the year fixed effect that controls for macroeconomic shocks, and ε_{it} is the structural error term related to the intangible share parameter μ_{it}.

We estimate equation (13) using the panel of US public firms described in Section 2.1, with σ as the parameter of interest. Since J reflects endogenous multinational expansion decisions that may covary with unobserved determinants of intangible intensity, OLS estimates of σ will be biased. We therefore instrument for $\log J$ 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 J_{it}. A pre-trends test ($p > 0.34$) provides evidence against differential pre-treatment trends between industries with different exposures, supporting the exclusion restriction[3].

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} 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 4 shows the estimation result of equation (13). The OLS coefficient on $\log J$, 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 $\log J$, 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 (2024b) estimate elasticities of substitution

[3]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.

Table 4: IV Estimates of the Elasticity of Substitution (σ)

	(1) OLS	(2) EU Bartik	(3) BIT Network	(4) Both
$\log J$	-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
Observations	148,682	148,682	148,682	148,682
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 $\log J$. 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: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

between software/ICT or IP capital and labor that exceed one.[4] 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.

6 Quantitative Results

This section presents preliminary quantitative results from the dynamic general equilibrium model developed in Section 4. The objective is not to provide a full quantitative decomposition of historical trends, but rather to illustrate the model's core mechanisms and discipline their magnitude. We proceed in three steps. First, we inspect firm-level policy functions and the implied selection patterns to clarify how productivity and intangible capital jointly govern multinational expansion. Second, we study comparative statics with respect to key technological and policy parameters—intangible investment costs, expansion costs, and barriers to foreign direct investment—to understand how different shocks operate through intensive and extensive margins. Third, we quantify the complementarity between intangible investment and multinational expansion using

[4] Lashkari, Bauer, and Boussard (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).

a set of transparent counterfactual experiments. Together, these results demonstrate how nonrival intangible capital shapes firm behavior and generates amplification effects that are central to the model's quantitative implications.

6.1 Inspecting the Mechanism

Multinational Expansion. Figure 5 provides an economic interpretation of the firm's expansion decision at the extensive margin. Foreign expansion entails a fixed cost, but allows the firm to deploy its intangible capital across multiple markets. Intangible capital therefore plays a central role because it is largely non-rival across locations: once developed, it can be leveraged abroad without proportional increases in production costs. As a result, firms with higher intangible capital face a higher marginal benefit from expansion, which explains the steep increase in expansion probabilities along the z dimension.

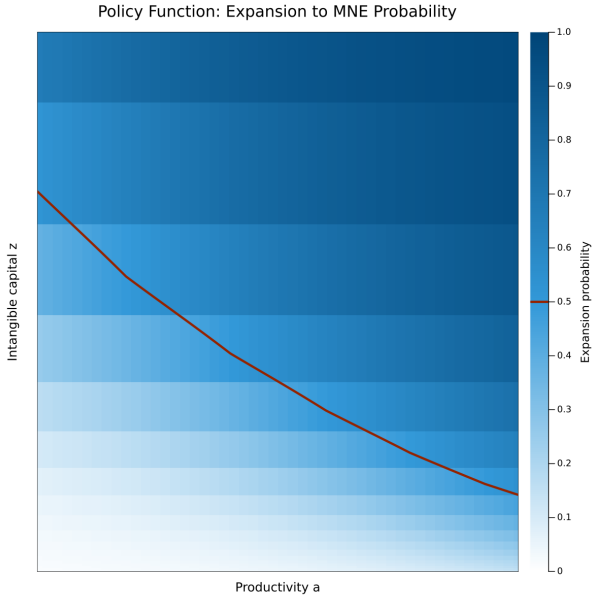


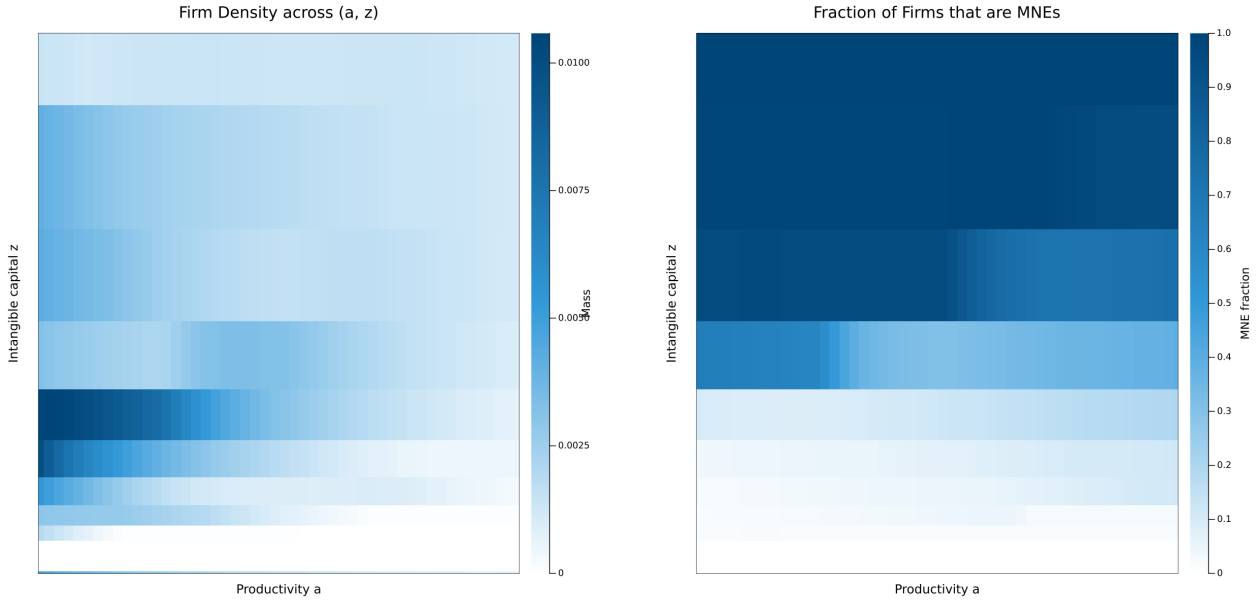
Figure 5: Policy function: expansion (to MNE) probability at $n = 0$.

Productivity affects the expansion decision through scale. More productive firms operate at a larger scale in the domestic market, which increases revenues and allows fixed expansion costs to be amortized over a larger output base. This mechanism raises the expansion probability in productivity, though more gradually than for intangible capital. The downward-sloping contour captures the substitutability between these two margins: high productivity can partially compensate for lower intangible capital, and vice versa, in making foreign expansion profitable.

Taken together, the figure highlights that multinational selection is not driven by productivity

alone. Instead, expansion reflects a joint interaction between scale (productivity) and scalability (intangible capital), with intangibles governing how effectively firms can translate size into profitable foreign presence.

Selection. Figure 6 illustrates how the expansion policy in Figure 5 maps into selection and composition in the stationary distribution. Panel (a) shows the joint density of firms over productivity a and intangible capital z. Firm mass is concentrated at relatively low levels of intangible capital, with productivity dispersion within each z bin. This reflects the fact that intangible accumulation is costly and that only a subset of firms finds it optimal to operate with high intangible intensity.



(a) Firm density over (a, z).

(b) Fraction of firms that are MNEs over (a, z).

Figure 6: Selection and composition in the stationary distribution.

Panel (b) reports, for each point in the (a, z) space, the fraction of firms that operate as multinationals. Multinational activity is negligible at low levels of intangible capital, regardless of productivity, but rises sharply with z. Conditional on sufficiently high intangible capital, even moderately productive firms exhibit a high probability of foreign expansion, while productivity alone is insufficient to generate multinational status in the absence of intangibles.

Taken together, the two panels highlight the nature of selection at the extensive margin. Although most firms are small and intangible-poor, multinational firms are disproportionately drawn from the upper tail of the intangible distribution. Productivity affects selection primarily by shifting firms along the scale dimension, but intangible capital determines whether scale can be effectively

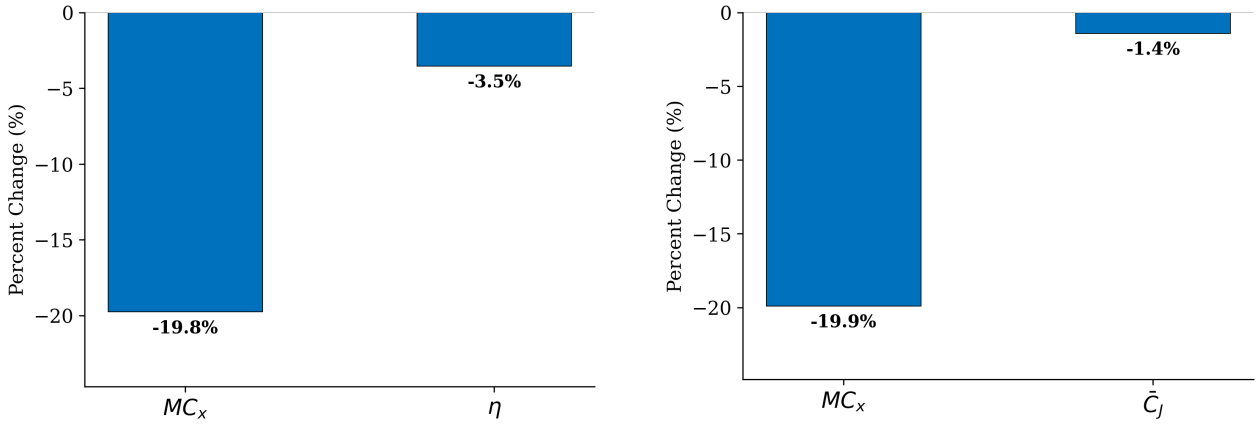
leveraged across markets. As a result, multinational firms are not simply the most productive firms in the economy; they are firms that combine sufficient productivity with scalable intangible assets, leading to a sharp sorting pattern in the stationary equilibrium.

The selection patterns in Figure 6 are consistent with several well-documented empirical facts about multinational firms. Empirically, multinational firms are larger, more productive, and substantially more intangible-intensive than purely domestic firms, but productivity alone does not fully explain multinational status. The model reproduces this feature by generating selection that is sharp in intangible capital but gradual in productivity. In particular, the absence of a strict productivity cutoff mirrors the empirical observation that many highly productive firms remain domestic, while some moderately productive firms operate internationally. In the model, this occurs because foreign expansion is primarily driven by the presence of scalable intangible assets, which determine whether scale advantages can be profitably extended across borders. As a result, multinational firms emerge as a selected subset of firms that combine sufficient productivity with high intangible intensity, rather than simply occupying the extreme upper tail of the productivity distribution.

6.2 Counterfactual Decomposition of Intangible Costs and Globalization Frictions

We next use the calibrated model to decompose the joint rise in multinational activity and intangible intensity into changes in intangible investment costs and changes in globalization frictions. The exercise compares an initial (1995) and a terminal (2023) period calibration. Starting from the terminal-period calibration, we vary the marginal cost of intangible investment, MC_x, together with either the fixed cost of multinational expansion, $\bar{C}_J$, or the variable FDI barrier, η, until the model matches the change in the targeted moments between the two periods. We then ask how much of the implied change in each target moment is accounted for by each parameter movement. This decomposition is useful because the model features strong cross-effects: lower intangible costs affect expansion incentives, and lower globalization frictions affect the return to intangible accumulation.

Figure 7 reports the parameter changes needed to rationalize the shift in the target moments between the initial and terminal periods. In both versions of the exercise, the dominant movement is a decline of roughly 20 percent in MC_x. By contrast, the globalization parameters move much less in percentage terms: the implied decline is about 3.5 percent for η and about 1.4 percent for $\bar{C}_J$. The first message of the counterfactual is therefore that the model attributes the bulk of the underlying structural change to cheaper intangible accumulation, with globalization frictions



(a) Decomposition using MC_x and η.

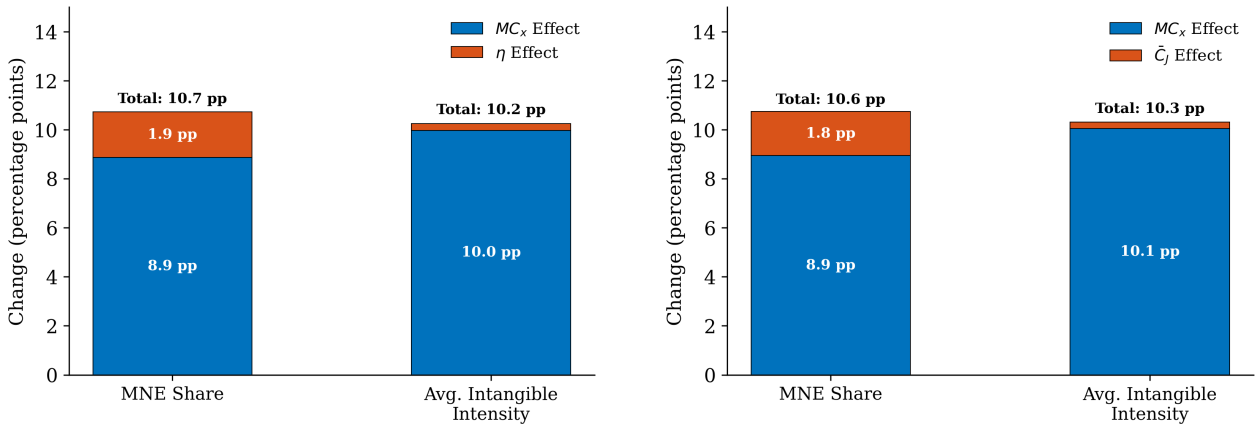
(b) Decomposition using MC_x and $\bar{C}_J$.

Figure 7: Parameter changes required to match the observed change in target moments.

adjusting by a smaller amount.

The more informative object is not the parameter change itself, but how those changes map into the targeted moments. Figure 8 reports this decomposition. In the baseline specification, the model attributes a 10.1 percentage point increase in the MNE share to a 3.2 percentage point contribution from MC_x and a 7.0 percentage point contribution from η. By contrast, the 7.7 percentage point rise in average intangible intensity is accounted for almost entirely by MC_x, which contributes 7.0 percentage points, while η contributes only 0.7 percentage points. This asymmetric pattern is economically intuitive. Lower intangible costs directly raise firms' desired stock of non-rival capital and therefore account for most of the rise in intangible intensity. Lower FDI frictions, in contrast, primarily affect the profitability of foreign operation and therefore account for a larger share of the increase in multinational participation.

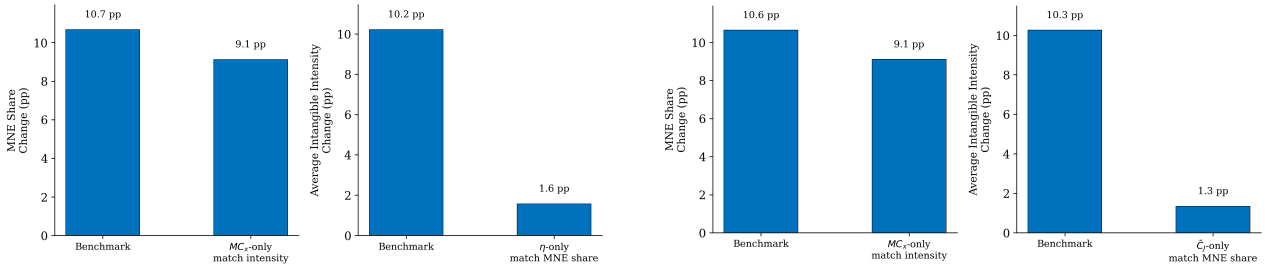
Panels (a) and (b) of Figure 8 show that this conclusion is robust, but they also reveal that the spillovers from intangible costs to multinational expansion are quantitatively large. In the decomposition based on (MC_x, η), the rise in the MNE share is 10.7 percentage points, of which 8.9 percentage points are attributed to MC_x and 1.9 percentage points to η. In the decomposition based on $(MC_x, \bar{C}_J)$, the corresponding contributions are 8.9 and 1.8 percentage points. For average intangible intensity, the result is even sharper: almost the entire increase is attributed to the decline in MC_x, with the contribution of η or $\bar{C}_J$ close to zero. The broader implication is that, once general-equilibrium and selection effects are taken seriously, lower intangible costs affect not only the intensive margin of investment but also the extensive margin of multinational expansion.



(a) Decomposition using η.

(b) Decomposition using $\bar{C}_J$.

Figure 8: Decomposition of changes in target moments into the effects of changes in MC_x and globalization frictions.



(a) Spillovers in the (MC_x, η) decomposition.

(b) Spillovers in the $(MC_x, \bar{C}_J)$ decomposition.

Figure 9: Single-parameter spillovers A target moments.

This finding follows naturally from the mechanism emphasized in the theory. A decline in MC_x raises the private return to intangible accumulation. Because intangible capital is non-rival, the same increase in z simultaneously raises the gains from serving foreign markets, so cheaper intangibles induce more firms to expand abroad. By contrast, a reduction in η or $\bar{C}_J$ mainly changes the profitability threshold for multinational operation. It therefore has a large effect on entry into multinational status, but a much weaker effect on firms' desired stock of intangibles conditional on operation.

Figure 9 isolates these cross-effects more directly. The left bars ask how much of the rise in the MNE share is generated when MC_x alone is moved to match the rise in intangible intensity. The answer is substantial: in both decompositions, changing MC_x alone produces a 9.1 percentage point increase in the MNE share, close to the benchmark increases of 10.7 and 10.6 percentage

points. The right bars ask the converse question: how much of the rise in intangible intensity is generated when η or $\bar{C}_J$ alone is moved to match the rise in the MNE share. Here the answer is much smaller. Matching the MNE share with η alone raises average intangible intensity by only 1.6 percentage points, and matching it with $\bar{C}_J$ alone raises average intangible intensity by only 1.3 percentage points, both far below the benchmark increases of 10.2 and 10.3 percentage points.

Taken together, these counterfactuals point to a clear asymmetry. Lower intangible investment costs generate large spillovers onto multinational expansion, while lower globalization frictions generate only modest spillovers onto intangible deepening. The decomposition therefore provides quantitative support for the paper's central mechanism: the rise of scalable intangible capital is a first-order driver of multinational expansion, and globalization amplifies that force primarily by raising the return to deploying intangibles across markets rather than by independently generating large increases in intangible investment.

7 Implications for Aggregate Productivity Growth

A central feature of our framework is that intangible capital is a nonrival factor of production: the same stock of intangibles developed at a firm's headquarters is deployed in full across each of its production locations. This property of the input, combined with heterogeneity in firm productivity, generates a distribution of intangible intensities across plants. A property of the production technology then determines whether this heterogeneity has implications beyond factor proportions: when the elasticity of substitution between intangible and rival inputs exceeds one, differences in intangible intensity across plants translate into differences in output elasticities. This section develops the implications of this elasticity heterogeneity for how changes in intangible investment and multinational expansion map into measured aggregate productivity.

We proceed with a bottom up approach, defining the total factor productivity at the plant level and aggregating up to the aggregate TFP levels, which we then convert to the growth rates. This is in contrast to the approach take Bureau of Economic Analysis (BEA) and Bureau of Labor Statistics (BLS), whose estimates rely on the top-down, aggregate input and output measures.

7.1 Nonrivalry, Intangible Intensity, and Output Elasticities

Consider a plant operated by a firm with intangible capital z, using rival inputs $r = k^{\alpha}\ell^{1-\alpha}$ hired locally. The plant's output is given by the CES production function

$$y = a\left[\mu z^{\frac{\sigma-1}{\sigma}} + (1-\mu)r^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\theta\sigma}{\sigma-1}}, \tag{14}$$

where σ is the elasticity of substitution between intangible and rival inputs. The output elasticity with respect to intangible capital is

$$\varepsilon_{y,z} = \theta \cdot \frac{\mu z^{\frac{\sigma-1}{\sigma}}}{\mu z^{\frac{\sigma-1}{\sigma}} + (1-\mu) r^{\frac{\sigma-1}{\sigma}}}. \tag{15}$$

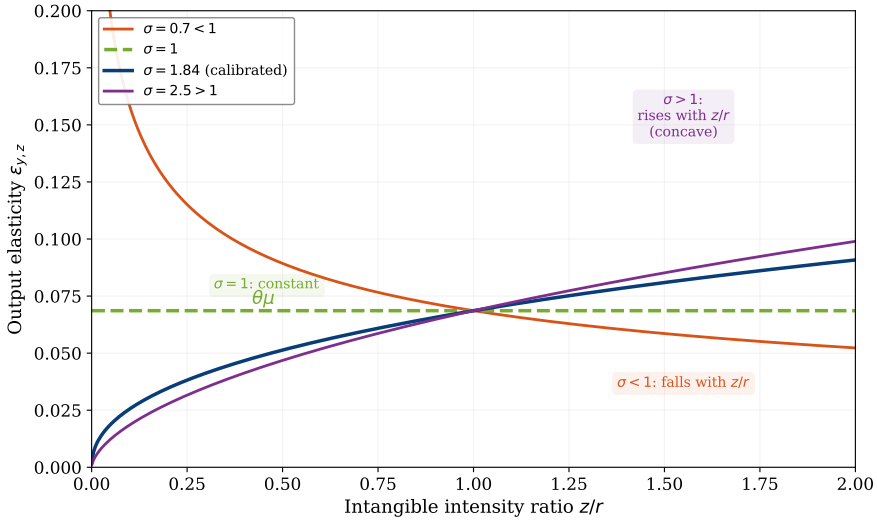
This elasticity depends on the plant's intangible intensity z/r whenever $\sigma \neq 1$. When $\sigma = 1$, the production function is Cobb-Douglas and the elasticity reduces to $\varepsilon_{y,z} = \theta\mu$ for all plants, regardless of their input mix.

In our framework, plants differ in their intangible intensity z/r as a consequence of two features: the nonrivalry of intangible capital and heterogeneity in firm productivity. Intangible capital is nonrival in the sense that a firm's entire stock z is deployed in each of its production locations without diminishing its productivity at any other location. A firm with productivity a that operates J plants hires rival inputs r_j locally in each location j, but deploys the same z in all of them. As shown in Section 3, higher-productivity firms optimally choose both greater intangible capital z and broader multinational scope J. Because z is nonrival, expanding to additional locations does not dilute the intangible input available to any individual plant. The net effect is that high-productivity multinational plants operate at systematically higher intangible intensity z/r than plants of domestic firms.

Whether this heterogeneity in intangible intensity has implications for output elasticities depends on the elasticity of substitution σ. When $\sigma = 1$, the output elasticity $\varepsilon_{y,z} = \theta\mu$ is constant across all plants: differences in z/r affect factor proportions but not the share of output attributable to intangible capital. When $\sigma > 1$, however, the output elasticity $\varepsilon_{y,z}$ is strictly increasing in z/r. Plants with higher intangible intensity—which, as established above, are disproportionately plants of high-productivity multinational firms—have a larger share of output attributable to intangible capital. The distribution of intangible intensities generated by nonrivalry and productivity heterogeneity therefore maps into a distribution of output elasticities across plants. Figure 10 illustrates this relationship for different values of σ, including our calibrated value of $\sigma = 1.84$.

To summarize, the nonrival nature of intangible capital, combined with heterogeneity in firm productivity, endogenously generates variation in plant-level intangible intensity. Under an elasticity of substitution greater than one, this variation produces heterogeneous output elasticities across plants, linked systematically to firms' productivity and multinational scope. The next subsection defines a sequence of plant-level productivity measures that progressively account for this elasticity heterogeneity, and quantifies the extent to which it matters for measured productivity

Figure 10: Output elasticity as a function of σ



growth.

7.2 Plant-Level Productivity Measures

We define four measures of plant-level total factor productivity that differ in their treatment of intangible capital and output elasticities. Because output markets are perfectly competitive in our framework, plant-level revenue productivity and quantity productivity differ only by a common output price. We adopt the revenue-product convention throughout to align with standard empirical practice.

Consider a plant operated in location j by a firm headquartered in country i with productivity a, intangible capital z, and n foreign affiliates. Let y_{ij} denote the plant's output, $r_{ij} = k_{ij}^{\alpha}\ell_{ij}^{1-\alpha}$ its rival input composite, and z_i the firm's intangible capital deployed in the plant.

Measure 1: Solow residual. The standard approach combines intangible and tangible capital into a single capital aggregate and imposes common output elasticities across all plants:

$$\ln tfp_{ij}^{1} = \ln y_{ij} - \alpha_k \ln(z_i + k_{ij}) - (1 - \alpha_k) \ln \ell_{ij}, \tag{16}$$

where α_k is a common capital share. This measure abstracts from the elasticity heterogeneity established in Section 6 in two ways. First, it imposes a uniform output elasticity on capital across all plants. Second, by aggregating intangible and tangible capital into a single input, it implicitly treats intangible capital as a rival factor—allocated to and used up at each plant individually—rather

than as a nonrival input deployed simultaneously across all of the firm's production locations.

Measure 2: Endogenous elasticities, intangible capital omitted. This measure allows the output elasticity on rival inputs to vary across plants according to the model-implied values, but does not account for intangible capital as a separate input:

$$\ln tfp_{ij}^{2} = \ln y_{ij} - \varepsilon_{y,r}^{ij} \ln r_{ij}, \tag{17}$$

where $\varepsilon_{y,r}^{ij}$ is the plant-specific output elasticity with respect to the rival composite, evaluated at the plant's equilibrium input mix. By using the correct rival-input elasticity, this measure removes the bias from imposing common factor shares. However, because intangible capital does not appear as a measured input, changes in z that raise output are absorbed entirely into the residual.

Measure 3: Endogenous elasticities, intangible capital included. This measure accounts for intangible capital as a separate input and uses plant-specific output elasticities for both intangible and rival factors:

$$\ln tfp_{ij}^{3} = \ln y_{ij} - \varepsilon_{y,z}^{ij} \ln z_i - \varepsilon_{y,r}^{ij} \ln r_{ij}, \tag{18}$$

where $\varepsilon_{y,z}^{ij}$ and $\varepsilon_{y,r}^{ij}$ are the plant-specific output elasticities with respect to intangible capital and the rival composite, respectively. This measure uses the correct inputs and the correct elasticities, but approximates the CES production function with a log-linear specification. This approximation is exact when $\sigma = 1$ but introduces an error for $\sigma \neq 1$, as the true technology features a nonlinear interaction between z and r that a log-linear form does not capture.

Measure 4: Production-function consistent. This measure computes the residual directly from the CES production function:

$$\ln tfp_{ij}^{4} = \ln y_{ij} - \frac{\theta\sigma}{\sigma - 1} \ln \left[\mu z_i^{\frac{\sigma-1}{\sigma}} + (1-\mu) r_{ij}^{\frac{\sigma-1}{\sigma}} \right]. \tag{19}$$

This is the model-consistent measure: it uses the true functional form, accounts for intangible capital as a nonrival input, and allows output elasticities to vary endogenously with the plant's input mix. Under the model's production technology, $tfp_{ij}^{4} = \ln a_i$ for domestic plants and $tfp_{ij}^{4} = \ln a_i + \ln(1-\eta)$ for plants of foreign multinationals, recovering the primitive productivity up to the FDI barrier.

The difference between consecutive measures isolates a distinct source of divergence. The gap $\ln tfp_{ij}^{1} - \ln tfp_{ij}^{2}$ captures the effect of imposing common rather than plant-specific output

elasticities, which we refer to as the *elasticity wedge*. The gap $\ln tfp^2_{ij} - \ln tfp^3_{ij}$ captures the effect of omitting intangible capital as a measured input, the *missing-z wedge*. The gap $\ln tfp^3_{ij} - \ln tfp^4_{ij}$ captures the error from approximating the CES technology with a log-linear specification, the *functional-form wedge*. Because aggregation uses common Domar weights across all four measures, this plant-level decomposition passes through to aggregate productivity: the total gap between aggregate measures 1 and 4 equals the sum of the three aggregate wedges.

7.3 Aggregation

We aggregate plant-level productivity measures to the economy level following the Domar-weighted approach of Hulten (1978). Let $s_{ij}(a, z, n)$ denote the share of plant (i, j) in aggregate output, where i indexes the firm's headquarters country, j the production location, and (a, z, n) the firm's state vector. Aggregate TFP under measure $k \in \{1, 2, 3, 4\}$ is

$$\ln TFP^k = \sum_j \int s_{ij}(a, z, n)\ \ln tfp^k_{ij}(a, z, n)\, d\Psi_i(a, z, n), \tag{20}$$

where $\Psi_i(a, z, n)$ is the stationary distribution of firms headquartered in country i over productivity, intangible capital, and multinational scope.

We focus on production taking place in a single country—the United States—and decompose aggregate TFP into contributions from three groups of plants distinguished by ownership and multinational status:

$$\ln TFP^k_{US} = \Omega_{NM} \ln TFP^k_{NM} + \Omega_{USMNE} \ln TFP^k_{USMNE} + \Omega_{FMNE} \ln TFP^k_{FMNE}, \tag{21}$$

where Ω_g is group g's share of U.S. output. The first group consists of plants operated by U.S. firms with no foreign affiliates ($n = 0$). The second consists of U.S. plants operated by U.S.-headquartered multinational firms ($n \geq 1$). The third consists of U.S. plants operated by foreign-headquartered multinational firms.

Computing changes across steady states. The quantitative exercises in Section 6 compare two steady-state equilibria, indexed by $t \in \{0, 1\}$. In each steady state, the distribution of firms, equilibrium prices, and optimal input choices differ. We compute aggregate TFP under each measure as the Domar-weighted sum of plant-level productivities in each steady state, and define the

aggregate change as the difference:

$$\Delta \ln TFP^k_{US} = \ln TFP^k_{US,1} - \ln TFP^k_{US,0}. \tag{22}$$

The content of this change differs across measures. For measure 4, plant-level TFP equals the firm's primitive productivity $\ln a_i$ (up to the FDI barrier for foreign multinational plants) and is therefore independent of the equilibrium. The entire change in aggregate TFP^4 reflects reallocation: shifts in the distribution of firms and in output shares across plants. For measures 2 and 3, plant-level TFP at a given state (a, z, n) varies across steady states because equilibrium prices alter the plant's optimal input mix, which in turn changes the plant-specific elasticities and the residual. The aggregate change therefore reflects both within-plant changes in measured productivity and reallocation across the firm distribution. For measure 1, an additional source of variation arises if the common capital share α_k is allowed to adjust across steady states to match aggregate cost shares in each equilibrium.

7.4 The Aggregate-Residual Approach

The Bureau of Economic Analysis and the Bureau of Labor Statistics compute aggregate productivity growth directly from aggregate quantities:

$$\Delta \ln TFP^{AR} = \Delta \ln Y - \sum_{j \in \mathcal{I}} \bar{s}_j \, \Delta \ln X_j, \tag{23}$$

where Y is aggregate output, X_j are aggregate inputs, and $\bar{s}_j = \frac{1}{2}(s_{j,0} + s_{j,1})$ is the Törnqvist-averaged cost share of input j across the two periods. We apply this aggregate-residual approach to model-generated data in order to replicate what the BEA and BLS would compute if they observed our model economy. We use the Törnqvist index because that is what the BEA and BLS use, and the point of the exercise is to compare their methodology against the bottom-up measures where we exploit full knowledge of the model. The Törnqvist convention is not a methodological choice we are making—it is a feature of the measurement technology we are evaluating. Details on the mapping between the model and the BEA/BLS methodology are provided in Appendix D.2.

When we replicate the aggregate-residual approach inside the model, we construct aggregate inputs as they would appear in national accounts. Intangible capital is booked in the headquarters country of the firm that developed it. The aggregate capital stock entering the BEA/BLS computation is therefore

$$K^{AR}_{US} = Z^{booked}_{US} + K^{tangible}_{US}, \tag{24}$$

where $K_{US}^{tangible}$ is total tangible capital employed in the United States and

$$Z_{US}^{booked} = \int z_i \, d\Psi_{US}(a, z, n) \tag{25}$$

is the total intangible capital of U.S.-headquartered firms. This differs from the intangible capital actually used in U.S. production,

$$Z_{US}^{used} = \int z_i \, d\Psi_{US}(a, z, n) + \sum_{j \neq US} \int z_i \, d\Psi_{j,US}(a, z, n), \tag{26}$$

which additionally includes the intangible capital deployed by foreign-headquartered multinationals in their U.S. affiliates. Because intangible capital is nonrival, these firms deploy their full stock z_i in each production location, including the United States, while national accounts record this capital in the firm's home country. The gap $Z_{US}^{used} - Z_{US}^{booked}$ is therefore driven entirely by inward multinational production: the aggregate-residual approach undercounts the intangible input used in U.S. production, attributing the resulting output to productivity rather than to factor accumulation.

7.5 Decomposition of Aggregate TFP Changes

We apply the four bottom-up measures and the aggregate-residual approach to the steady-state equilibria computed in Section 6.1. Table 5 reports the change in aggregate TFP under each measure between the pre- and post-period steady states.

The top panel of Table 5 reveals that the choice of productivity measure has first-order consequences for the measured change in aggregate TFP. The Solow residual (measure 1) implies a decline of 4.70 percentage points, while the production-function-consistent measure (measure 4) implies a decline of only 0.24 percentage points. The gap between these two measures—4.46 percentage points—represents the total bottom-up measurement wedge: the extent to which standard growth accounting, applied at the plant level with common factor shares, mismeasures the productivity consequences of the intangible revolution and multinational expansion.

This total wedge is dominated by the elasticity wedge (−4.35 percentage points), which captures the bias from imposing common output elasticities across all plants. Intuitively, the Solow residual assigns a uniform capital share to every plant, but plants of intangible-intensive multinational firms have substantially higher output elasticities with respect to intangible capital than plants of domestic firms. When the intangible revolution shifts activity toward these high-elasticity

Table 5: Aggregate TFP Growth Under Alternative Measures

	$\Delta \ln TFP^k$ (%)
(1) Solow residual	-4.70
(2) Endogenous elasticities, z omitted	-0.35
(3) Endogenous elasticities, z included	+2.04
(4) Production-function consistent	-0.24
(AR) Aggregate residual (BEA/BLS)	+1.19
Measurement wedges (bottom-up)	
$\Delta \ln TFP^1 - \Delta \ln TFP^2$ (elasticity wedge)	-4.35
$\Delta \ln TFP^2 - \Delta \ln TFP^3$ (missing-z wedge)	-2.39
$\Delta \ln TFP^3 - \Delta \ln TFP^4$ (functional-form wedge)	+2.28
$\Delta \ln TFP^1 - \Delta \ln TFP^4$ (**total bottom-up wedge**)	-4.46
Aggregate-residual wedge	
$\Delta \ln TFP^{AR} - \Delta \ln TFP^4$ (**total AR wedge**)	+1.44

Notes: This table reports the change in aggregate TFP between the pre- and post-period steady states under four bottom-up measures and the aggregate-residual approach. Bottom-up measures are computed as Domar-weighted sums of plant-level productivities in each steady state. The aggregate-residual measure replicates the BEA/BLS methodology applied to model-generated data, using intangible capital booked in the headquarters country as recorded in national accounts. Measurement wedges are defined as differences between consecutive measures. All entries are in percentage points.

plants, the Solow residual systematically undercounts the contribution of intangible capital deepening to output growth, generating a large negative residual. The magnitude of this wedge reflects the calibrated elasticity of substitution $\sigma = 1.84$: because intangible and rival inputs are gross substitutes, reductions in the cost of intangible investment induce firms to raise their intangible intensity, which in turn raises the output elasticity of intangible capital at plants that were already intangible-intensive. Imposing a common capital share absorbs none of this heterogeneity, and the resulting measurement error is amplified by the reallocation of output toward multinational firms.

The missing-z wedge (–2.39 percentage points) captures the effect of omitting intangible capital as a measured input. When plant-specific elasticities are used but intangible capital is excluded from the input set, increases in z that raise output are absorbed entirely into the productivity residual. The magnitude of this wedge reflects the substantial growth in intangible capital across both steady states: the roughly 20 percent decline in the marginal cost of intangible investment induces firms to accumulate significantly more intangible capital, and when this accumulation is not accounted for, it appears as productivity growth rather than factor deepening. The functional-form wedge (+2.28 percentage points) partially offsets the missing-z wedge. This term captures the ap-

proximation error from using a log-linear specification to account for intangible capital when the true technology is CES with $\sigma \neq 1$. The positive sign reflects the fact that the log-linear approximation overestimates the marginal product of intangible capital at high intangible intensities when $\sigma > 1$, thereby over-subtracting the contribution of z and understating the residual relative to the true CES form.

The aggregate-residual measure, which replicates the BEA/BLS methodology applied to model-generated data, implies an increase of 1.19 percentage points—qualitatively and quantitatively different from the production-function-consistent measure, which shows a decline of 0.24 percentage points. The total aggregate-residual wedge of +1.44 percentage points reflects two distinct and roughly equally important channels. The first is a common-share channel (+0.71 percentage points): the aggregate-residual approach uses aggregate cost shares as proxies for output elasticities, and the Törnqvist aggregation partially smooths over the heterogeneity in plant-level elasticities, generating a positive bias relative to the production-function-consistent measure. The second is a booking channel (+0.73 percentage points), specific to the open-economy setting: national accounts record intangible capital in the headquarters country of the firm that developed it, regardless of where that capital is deployed in production. Because intangible capital is nonrival, foreign-headquartered multinationals deploy their full stock z_i in their U.S. affiliates, but this capital is booked abroad. The resulting gap between Z_{US}^{used} and Z_{US}^{booked} means that the aggregate-residual approach undercounts the intangible input actually used in U.S. production, attributing the associated output to productivity rather than to factor accumulation. As multinational activity expands, the booking wedge widens: the gap between used and booked intangible capital nearly doubles between the pre- and post-period steady states, driven by the growth in inward multinational production. Taken together, the two channels imply that the BEA/BLS methodology would record positive TFP growth of over one percentage point in a model economy where the true change in production-function-consistent productivity is essentially zero.

To further disentangle the sources of aggregate TFP changes under the production-function-consistent measure, we decompose aggregate TFP in each steady state into an unweighted mean and a covariance term following Olley and Pakes (1996):

$$\ln TFP_{US,t}^{4} = \overline{\ln tfp^{4}}_{t} + \sum_{i} \left(s_{i,t} - \bar{s}_{t}\right) \left(\ln tfp_{i,t}^{4} - \overline{\ln tfp^{4}}_{t}\right), \tag{27}$$

where $\overline{\ln tfp^{4}}_{t}$ is the unweighted mean of plant-level log productivity under measure 4 and $s_{i,t}$ is the plant's output share. The first term captures the average level of production-function-consistent

Table 6: Olley–Pakes Decomposition of Production-Function-Consistent TFP Growth

	Total	Mean	Allocative Efficiency	
		(Composition)	Within	Between
Non-multinationals	+0.37	+1.04	-0.79	+0.12
U.S. MNEs (domestic plants)	-0.35	-0.54	+0.28	-0.09
Foreign MNEs (U.S. plants)	-0.27	-0.45	+0.24	-0.06
Aggregate	-0.24	+0.04	-0.26	-0.03

Notes: This table reports the Olley–Pakes decomposition of the change in aggregate production-function-consistent TFP (measure 4) between the pre- and post-period steady states. Plant-level log productivities are normalized by subtracting the pre-period aggregate level before computing group contributions, so that the Total column reflects each group's contribution to TFP growth rather than levels. The mean component captures changes in the unweighted average of plant-level log productivity. The allocative efficiency component captures changes in the covariance between productivity and output share, decomposed into within-group and between-group reallocation across three plant groups. All entries are in percentage points. Rows sum to the Total column; columns sum to the Aggregate row.

productivity across plants (composition). The second term captures allocative efficiency: the covariance between plant-level productivity and market share. We further decompose the covariance term into within-group and between-group components across the three plant groups defined in Section 7.3.

Table 6 decomposes the change in production-function-consistent TFP into composition and reallocation channels across three groups of plants operating in the United States. Because measure 4 recovers the primitive productivity $\ln a_i$ (up to the FDI barrier $\ln(1-\eta)$ for foreign multinational plants), this decomposition is not affected by the elasticity heterogeneity that drives the wedges in Table 5—it isolates the pure reallocation consequences of the intangible revolution and multinational expansion.

The aggregate change under measure 4 is small: −0.24 percentage points. This near-zero result reflects the fact that the counterfactual does not alter the distribution of firm-level productivities a—the Pareto distribution and its parameters are held fixed across steady states. The entire change in aggregate TFP^4 therefore reflects shifts in the composition of plants operating in the United States and in the allocation of output across them, rather than improvements in the underlying technology available to any individual firm.

The group-level decomposition reveals that this small aggregate decline masks modest but economically interpretable offsetting forces. Non-multinational plants contribute +0.37 percentage points to aggregate TFP growth, while U.S. MNE domestic plants and foreign MNE plants

contribute −0.35 and −0.27 percentage points, respectively. The magnitudes of these group-level contributions are comparable to the aggregate change itself, indicating that the decomposition captures genuine contributions to productivity growth rather than mechanical level effects.

The positive contribution from non-multinationals is driven by the mean (composition) component (+1.04 percentage points). As the MNE share rises from the pre- to the post-period equilibrium, firms that transition to multinational status exit the non-multinational group. These departing firms are among the more productive non-multinationals—consistent with the selection patterns documented in Section 2—and the resulting contraction of the group raises the output-weighted average productivity of the remaining non-multinational plants. This composition effect is partially offset by a decline in within-group allocative efficiency (−0.79 percentage points): as intangible costs fall, non-multinational firms with relatively higher intangible intensity expand their output shares, but these are not necessarily the highest-a firms within the group, reducing the covariance between productivity and market share. The between-group term for non-multinationals is small and positive (+0.12 percentage points), reflecting a modest shift of output toward the non-multinational group's relatively high average productivity level as the group contracts.

The negative contributions from the two multinational groups reflect the mirror image of this selection mechanism. For U.S. MNEs, the mean component is −0.54 percentage points: the entry of new multinationals dilutes the average productivity of the group, because the marginal firms that find expansion profitable are the most productive among former non-multinationals but less productive than incumbent multinationals. The within-group allocative efficiency term is positive (+0.28 percentage points), indicating that within the MNE group, output shares shift modestly toward higher-productivity firms as cheaper intangibles allow the most productive multinationals to scale further. Foreign MNE plants exhibit an analogous pattern: the mean component is −0.45 percentage points, reflecting the entry of marginal foreign multinationals whose U.S. affiliates additionally carry the productivity penalty $\ln(1-\eta)$ associated with the variable FDI barrier, while within-group reallocation is modestly positive (+0.24 percentage points). Between-group reallocation is essentially zero in aggregate (−0.03 percentage points), indicating that the shift of output from non-multinational to multinational plants does not, on net, move activity toward groups with systematically higher or lower production-function-consistent productivity.

The key implication of the joint reading of Tables 5 and 6 is that the intangible revolution and multinational expansion have large consequences for *measured* productivity but small consequences for *true* productivity as recovered by the model. Table 6 shows that production-function-consistent TFP is nearly unchanged (−0.24 percentage points), with the small decline reflecting

offsetting composition effects: the non-multinational group gains from the exit of marginal firms to MNE status, while both multinational groups are diluted by the entry of less productive firms at the extensive margin. Table 5 shows that the same forces generate a 4.46 percentage point bottom-up wedge between the Solow residual and the production-function-consistent measure, and a 1.44 percentage point wedge between the aggregate-residual approach and the true measure. The BEA/BLS methodology would record positive TFP growth exceeding one percentage point in a model economy where the underlying technology is essentially unchanged. This bias reflects two channels of roughly equal magnitude: the use of common factor shares that fail to account for heterogeneous output elasticities across plants, and the booking of intangible capital in the headquarters country rather than where it is deployed in production. Importantly, both channels are amplified by the same underlying forces—rising intangible intensity and expanding multinational production—implying that the divergence between measured and true productivity will widen as these trends continue.

8 Conclusion

This paper studies the joint rise of intangible capital and multinational production through the lens of non-rivality. Using firm-level data, we document that multinational firms are systematically more intangible-intensive than domestic firms, and that firms increase their intangible intensity following multinational entry. To interpret these patterns, we develop a quantitative model in which firms jointly choose intangible investment and global expansion. The central implication is that cheaper scalable intangibles are a primary force behind both intangible deepening and multinational expansion, while globalization raises the return to deploying those intangibles across markets. We find robust evidence that the elasticity of substitution between intangible and rival inputs exceeds one. As a result, firms optimally tilt their input mix toward intangible capital as their geographic scope expands, increasing intangible intensity. At the same time, this implies substantial cross-firm heterogeneity in the output elasticity of intangible capital.

These mechanisms have important implications for productivity measurement. In our framework, the rise of intangibles and multinational production generates divergence between production-function-consistent productivity and standard residual-based measures. The reason is twofold: first, common-share growth accounting does not accommodate the elasticity heterogeneity induced by non-rival intangibles; second, national accounts may record intangible capital where it is booked rather than where it is used in production. More broadly, the paper shows that understanding the aggregate consequences of the intangible economy requires treating intangible accumulation, firm organization, and productivity measurement as jointly determined outcomes.

References

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

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.

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.

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

Chaney, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," American Economic Review, 98, 1707–21.

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

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

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

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

Hulten, C. R. (1978): "Growth Accounting with Intermediate Inputs," The Review of Economic Studies, 45, 511–518.

Jerbashian, V. (2024a): "On the elasticity of substitution between labor and ict and ip capital and traditional capital," Working paper.

——— (2024b): "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.

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.

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

Levinsohn, J. and A. Petrin (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," Review of Economic Studies, 70, 317–341.

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.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," Econometrica, 71, 1695–1725.

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

Olley, G. S. and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," Econometrica, 64, 1263–1297.

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.

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.

Appendix

A Proofs of the Theory Section

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

$$\rho \equiv \frac{\sigma - 1}{\sigma} \in (0, 1), \qquad F(z, l) = \left[\mu z^{\rho} + (1 - \mu) l^{\rho}\right]^{\alpha/\rho}.$$

The firm's problem is

$$\max_{l,z,J} \left\{ J\left[p\, a F(z, l) - w l\right] - MC \cdot z - \bar{C}^J J^{\gamma} \right\}.$$

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

$$Jpa\alpha \left[\mu z^{\rho} + (1 - \mu) l^{\rho}\right]^{\frac{\alpha}{\rho} - 1} (1 - \mu) l^{\rho - 1} = w, \tag{28}$$

$$Jpa\alpha \left[\mu z^{\rho} + (1 - \mu) l^{\rho}\right]^{\frac{\alpha}{\rho} - 1} \mu z^{\rho - 1} = MC. \tag{29}$$

Dividing (29) by (28) yields

$$\frac{MC}{w} = \frac{\mu}{1 - \mu} \left(\frac{z}{l}\right)^{\rho - 1}, \qquad \text{so} \qquad \frac{z}{l} = \left(\frac{1 - \mu}{\mu} \frac{w}{MC}\right)^{\sigma}. \tag{30}$$

Because $s^z \equiv \frac{MC \cdot z}{Jwl + MC \cdot z}$, substituting z/l from (30) gives

$$s^z = \left[\left(\frac{1 - \mu}{\mu}\right)^{\sigma} \left(\frac{w}{MC}\right)^{1 - \sigma} J^{1 - \sigma} + 1\right]^{-1}, \tag{31}$$

which is equation (3). The first-order condition with respect to J is

$$\pi(a, l, z) = \gamma \bar{C}^J J^{\gamma - 1}, \tag{32}$$

hence

$$J = \left(\frac{\pi(a, l, z)}{\gamma \bar{C}^J}\right)^{\frac{1}{\gamma - 1}}, \tag{33}$$

which is equation (4). Equations (31) and (33) imply three sign restrictions used throughout:

$$\frac{\partial s^z}{\partial J} > 0, \qquad \frac{\partial J}{\partial \pi} > 0, \qquad \frac{\partial \pi(a, l, z)}{\partial a} > 0. \tag{34}$$

The first inequality follows from $\sigma > 1$, so $1 - \sigma < 0$ in (31); the second from $\gamma > 1$ in (33); and the third from $F(z, l) > 0$.

Proof of Proposition 1. We prove the two comparative statics for $\beta \in \{MC, \bar{C}^J\}$. Fix first J. From (31), a rise in MC lowers z/l in (30) and therefore lowers s^z, so $\partial s^z/\partial MC|_J < 0$. Likewise, because $1 - \sigma < 0$, (31) implies $\partial s^z/\partial J > 0$. Next, (33) gives $\partial J/\partial \bar{C}^J < 0$ immediately. Moreover, a rise in MC raises the direct cost of intangible capital, reduces the optimal z, lowers per-market profit $\pi(a, l, z)$, and therefore lowers J through (33). Hence $\partial J/\partial MC < 0$.

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

$$\frac{\partial s^z}{\partial \beta} = \left.\frac{\partial s^z}{\partial \beta}\right|_J + \frac{\partial s^z}{\partial J}\frac{\partial J}{\partial \beta}.$$

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

Proof of the Lemma. The decomposition

$$\frac{\partial s^z}{\partial \beta} = \left.\frac{\partial s^z}{\partial \beta}\right|_J + \frac{\partial s^z}{\partial J}\frac{\partial J}{\partial \beta}$$

is the chain rule applied to $s^z = s^z(J, \beta)$. It remains to sign the amplification term. Equation (31) implies $\partial s^z/\partial J > 0$ because $\sigma > 1$. Proposition 1 established that $\partial J/\partial \beta < 0$ for $\beta \in \{MC, \bar{C}^J\}$. Hence

$$\frac{\partial s^z}{\partial J}\frac{\partial J}{\partial \beta} < 0.$$

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

Proof of the Complementarity Proposition. Let $y \in \{s^z, J\}$. By definition,

$$\varepsilon_{y,(\bar{C}^J, MC)} = \varepsilon_{y,\bar{C}^J} + \varepsilon_{y,MC} + \frac{\partial^2 \ln y}{\partial \ln \bar{C}^J \partial \ln MC}.$$

It is therefore sufficient to show that the cross-partial term is negative. We first consider J. Taking logs in (33) gives

$$\ln J = \frac{1}{\gamma - 1}\left(\ln \pi(a, l, z) - \ln \gamma - \ln \bar{C}^J\right).$$

Hence

$$\frac{\partial \ln J}{\partial \ln MC} = \frac{1}{\gamma - 1} \frac{\partial \ln \pi}{\partial \ln MC} < 0.$$

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

$$\frac{\partial^2 \ln \pi}{\partial \ln \bar{C}^J \partial \ln MC} > 0 \qquad \Longrightarrow \qquad \frac{\partial^2 \ln J}{\partial \ln \bar{C}^J \partial \ln MC} < 0.$$

Thus $|\varepsilon_{J,(\bar{C}^J, MC)}| > |\varepsilon_{J,\bar{C}^J}| + |\varepsilon_{J,MC}|$.

Now consider s^z. Since s^z is increasing in J by (31), the same interaction carries over: a reduction in $\bar{C}^J$ raises J, and a larger J strengthens the effect of lower MC on the optimal intangible share. Therefore

$$\frac{\partial^2 \ln s^z}{\partial \ln \bar{C}^J \partial \ln MC} < 0,$$

which implies

$$|\varepsilon_{s^z,(\bar{C}^J, MC)}| > |\varepsilon_{s^z,\bar{C}^J}| + |\varepsilon_{s^z,MC}|.$$

Combining the two cases proves the proposition. □

Proof of Proposition 3. Equation (33) shows that J is strictly increasing in per-market profit $\pi(a, l, z)$. Since

$$\pi(a, l, z) = p\, aF(z, l) - wl,$$

we have $\partial \pi / \partial a = pF(z, l) > 0$ holding inputs fixed. Therefore higher-productivity firms choose larger J, so $\varepsilon_{J,a} > 0$.

Because (31) implies $\partial s^z / \partial J > 0$, any increase in productivity that raises J also raises the optimal intangible intensity. Hence $\varepsilon_{s^z,a} > 0$. This proves both claims. □

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

$$\frac{\partial^2 \ln J}{\partial \ln a \, \partial \ln MC} < 0.$$

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

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

Hence

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

An analogous argument applies to the pair $(a, \bar{C}^J)$. A more productive firm has higher profits, so it operates on a steeper part of the expansion schedule (33). A reduction in $\bar{C}^J$ therefore induces a larger increase in J for more productive firms, which in turn induces a larger rise in s^z. Thus

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

and therefore

$$|\varepsilon_{y,(a,\bar{C}^J)}| > |\varepsilon_{y,a}| + |\varepsilon_{y,\bar{C}^J}|.$$

This proves the proposition. □

B Measurement of Intangible Capital

This appendix provides a detailed discussion of our approach to measuring intangible capital at the firm level. We describe the components of balance-sheet intangibles under U.S. Generally Accepted Accounting Principles (GAAP), explain the construction of internally-generated knowledge capital from research and development expenditures, and justify our exclusion of organizational capital, capitalized from selling, general, and administrative (SG&A) expenses. We conclude by discussing the consistency of our measure with the existing intangible capital literature.

B.1 Balance-Sheet Intangible Capital Under GAAP

Balance-sheet total intangible capital (INTAN) in Compustat consists mostly of acquired intangible assets, which fall into the following four broad categories:

1. **Intellectual Property and Creative Assets:** Copyrights; patents; computer software patents; trademarks; subscription lists; design costs; engineering drawings; blueprints; software or software costs.

2. **Contractual, Legal, and Regulatory Rights:** Licenses; distribution rights and agreements; franchise and franchise fees; and similar contractual arrangements.

3. **Acquisition-Related Intangibles:** Goodwill; excess of cost or premium of acquisition; covenants not to compete; and other intangibles recognized in business combinations.

4. **Lease-Related Intangibles:** Leases and lease acquisition costs; leasehold expense; and related capitalized amounts.

In addition to acquired intangibles, GAAP permits, and in some cases requires, capitalization of certain internally-developed software costs. Per Statement of Financial Accounting Standards (SFAS) No. 86, issued in 1986, software costs (including software licenses, software technology, website development costs, and firmware) are required to be capitalized once technological feasibility is established.[5] Per Statement of Position (SOP) 98-1, first issued in 1998 with subsequent revisions, internal-use software is also to be capitalized during the application development stage.[6]

[5] SFAS No. 86 applies to software developed for sale, lease, or other marketing purposes. Costs incurred prior to establishing technological feasibility are expensed as R&D.

[6] SOP 98-1 was later codified into Accounting Standards Codification (ASC) 350-40 under the Financial Accounting Standards Board's codification project.

B.2 What Balance-Sheet Measures Miss

Despite the breadth of assets captured by balance-sheet intangibles, this measure omits two economically important cateogries of intangible investment that firms expense rather than capitalize under GAAP.

Internally-Generated Intellectual Property. U.S. GAAP generally prohibits the capitalization of internally-generated research and development expenditures. [7] As a result, balance-sheet intangibles exclude the cumulative stock of knowledge capital generated through in-house R&D activities, precisely the patents, proprietary technologies, and process innovations that are central to our theoretical framework.

To construct a measure of internally-generated knowledge capital, we follow the perpetual inventory method applied to firm-level data by Peters and Taylor (2017). Let $\text{R\&D}_{jt}$ denote R&D expenditure for firm j in year t, and let $\delta_n^{\text{R\&D}}$ denote an industry-specific depreciation rate for knowledge capital in industry n. We construct the stock of knowledge capital as:

$$K_{jt}^{\text{R\&D}} = (1 - \delta_{n(j)}^{\text{R\&D}}) K_{j,t-1}^{\text{R\&D}} + \text{R\&D}_{jt}, \tag{35}$$

with the initial stock set using the pre-sample growth rate method:

$$K_{j0}^{\text{R\&D}} = \frac{\text{R\&D}_{j0}}{g_n + \delta_n^{\text{R\&D}}}, \tag{36}$$

where g_n is the average R&D growth rate in industry n.

Following Peters and Taylor (2017), we use industry-specific depreciation rates ranging from 15% to 40%, reflecting heterogeneity in the pace of knowledge decay and obsolescence across sectors. Industries with rapid technological change (e.g., information technology, pharmaceuticals) receive higher depreciation rates, while industries with more durable knowledge stocks (e.g., chemicals, machinery) receive lower rates.

Organizational Capital. A second category of intangible investment that is expensed under GAAP is organizational capital. These are firm-specific assets embodied in business processes, human capital structures, brand equity, and supply chain relationships. Following Eisfeldt and Papanikolaou (2013), a common approach in the literature is to construct organizational capital by

[7] This contrasts with International Financial Reporting Standards (IFRS), which permit capitalization of development costs, once certain criteria are met.

capitalizing a fraction of selling, general, and administrative (SG&A) expenses:

$$K_{jt}^{\text{SG\&A}} = (1 - \delta^{\text{SG\&A}})K_{j,t-1}^{\text{SG\&A}} + \gamma \cdot \text{SG\&A}_{jt}, \tag{37}$$

where γ represents the share of SG&A expenditures that reflect investment in long-term organizational capital (typically set to 30%) and $\delta^{\text{SG\&A}}$ is the depreciation rate (typically $15-20\%$). The rationale is that SG&A expenses include salaries, consulting fees, marketing, and human resources expenditures, a portion of which builds durable organizational assets (e.g., marketing expenditures build brand equity, and training/HR expenses build human capital).

B.3 Our Measurement Approach

Our measure of intangible capital is the sum of balance-sheet intangibles and internally-generated knowledge capital from capitalized R&D:

$$\text{Intangible Capital}_{jt} = \text{INTAN}_{jt} + K_{jt}^{\text{R\&D}}. \tag{38}$$

We deliberately exclude organizational capital constructed from SG&A expenses, despite its use in related literature.

Theoretical Justification. Our exclusion of organizational capital is motivated by the theoretical framework developed in Section 3, which centers on the nonrivalry of intangible capital. Intangible assets are nonrival to the extent that they can be deployed simultaneously across multiple production locations without diminishing their productivity at any single site. This property generates the complementarity between intangible investment and multinational expansion that is central to our analysis: a firm that develops a patent, algorithm, or product design at headquarters can deploy it across all foreign subsidiaries at zero marginal cost.

Knowledge capital constructed from R&D expenditures captures the canonical nonrival inputs: patents, proprietary software, product designs, chemical formulas, and other forms of codified knowledge that can be scaled globally. These assets are precisely what multinational firms transfer across borders through licensing arrangements, and they constitute the primary vehicle for the profit-shifting mechanisms (transfer pricing of intellectual property) emphasized in our framework.

Organizational capital, by contrast, includes substantial location-specific and rival components alongside its nonrival elements. While organizational capital encompasses some scalable assets

(e.g., codifiable management practices, replicable business processes, and transferable brand equity), it also includes investments that are inherently tied to specific locations and cannot be deployed nonrivalrously across borders, such as local employee training, market-specific advertising, regional supplier relationships, and management attention.

By excluding organizational capital from our measure, we ensure that our empirical proxy corresponds closely to the theoretically relevant concept: intangible assets that can be deployed nonrivalrously across production locations and that drive the scalable complementarity between intangible investment and multinational expansion.

Empirical Justification. Our measurement choice also reflects the fact that the profit-shifting mechanisms central to our analysis (i.e., transfer pricing of intellectual property, cross-border licensing arrangements, and the reallocation of intangible-related income to low-tax jurisdictions) operate primarily through knowledge capital rather than organizational capital. Tax authorities and international guidelines (e.g., OECD Transfer Pricing Guidelines) focus on the valuation and allocation of returns to patents, trademarks, and similar intellectual property, not on the implicit value of organizational practices or human capital.

B.4 Within-Industry Growth Over Time

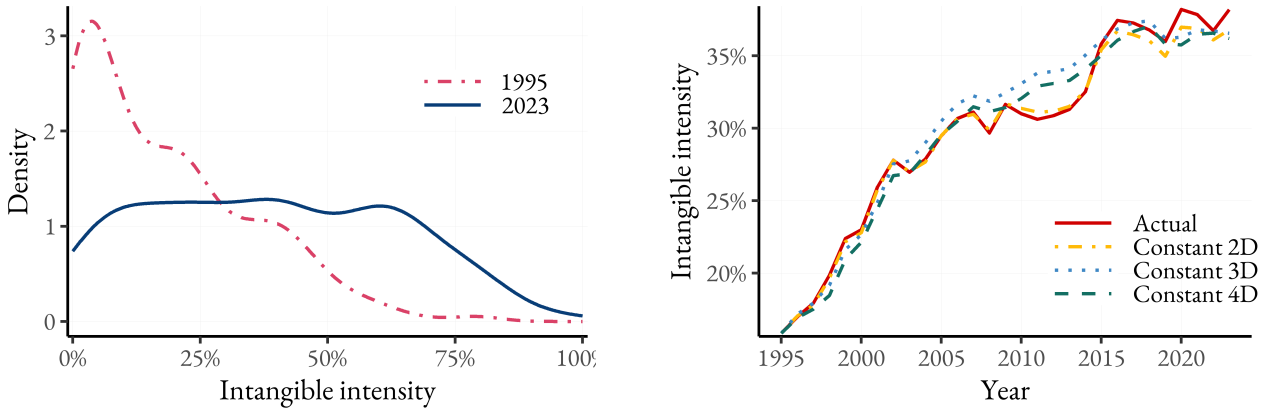
A natural concern is whether the aggregate trend documented in Figure 1 reflects within-industry increases in intangible intensity or compositional shifts toward high-intensity industries. Figure 11 addresses this concern: Panel (a) shows a broad-based rightward shift in the cross-industry distribution of intangible intensity between 1995 and 2023, while Panel (b) demonstrates that counterfactual series holding industry revenue shares fixed at 1995 levels track the actual series closely, confirming that the aggregate trend is driven overwhelmingly by within-industry increases.

B.5 Intangible Investment and FDI costs

The secular trends documented in Figure 1 occurred alongside substantial declines in the costs of intangible investment and foreign expansion. Figure 12 presents evidence on both margins.

Panel (a) presents the BEA Price Index for Private Fixed Investment in Intellectual Property Products (IPP), deflated by CPI and normalized to 100 in 2017. The overall IPP price index declined from approximately 150 in 1995 to 85 in 2023, representing a roughly 40% reduction in the real cost of intangible investment over this period. The decline is particularly pronounced for software, which fell from approximately 240 in 1995 to around 70 in 2023, a nearly 70% reduction.

Figure 11: Intangible Intensity Across Industries Over Time

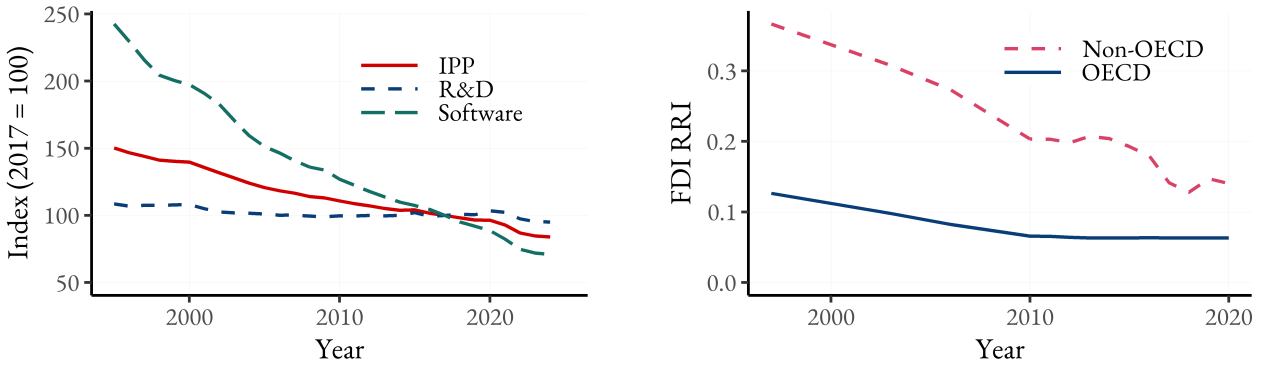


(a) Intangible Intensity Density in 1995 and 2023

(b) Intangible Intensity with Fixed Industry Shares

Notes: This figure examines whether the secular increase in intangible intensity reflects within-industry changes or compositional shifts across industries. Panel (a) presents kernel density estimates of intangible intensity aggregated to 4-digit NAICS industries using firm revenue weights, separately for 1995 (dashed line) and 2023 (solid line). The rightward shift indicates a broad-based increase in intangible intensity across the distribution rather than a reallocation toward high-intensity industries. Panel (b) displays counterfactual intangible intensity series constructed by holding industry revenue shares fixed at their 1995 levels while allowing within-industry intangible intensities to evolve. The "Constant 2D," "Constant 3D," and "Constant 4D" series fix industry composition at the 2-digit, 3-digit, and 4-digit NAICS level, respectively. The close tracking of counterfactual series with the actual series (solid red line) confirms that the aggregate trend is driven primarily by within-industry increases. For more granular industry definitions than 2 digits, we restrict attention to NAICS industries present throughout the 1995–2023 sample period.

Figure 12: Declining Costs of Intangible Investment and FDI Barriers: 1995-2023



(a) Intangible Investment Costs

(b) FDI Regulatory Restrictiveness

Notes: This figure documents the decline in the costs of intangible investment and foreign direct investment over the sample period. Panel (a) presents the Bureau of Economic Analysis (BEA) Price Index for Private Fixed Investment in Intellectual Property Products (IPP), deflated by the Consumer Price Index (CPI) and normalized to 100 in 2017. Three series are shown: the aggregate IPP index (solid red), the R&D component (dashed blue), and the software component (dashed green). The software price index exhibits the steepest decline, falling from approximately 240 in 1995 to 70 in 2023, reflecting rapid technological progress in information technology. Panel (b) displays the OECD FDI Regulatory Restrictiveness Index, which measures statutory restrictions on foreign direct investment including foreign equity limitations, screening and approval requirements, restrictions on key foreign personnel, and operational restrictions. The index ranges from 0 (fully open) to 1 (fully closed). Separate series are shown for OECD member countries (solid blue) and non-OECD countries (dashed red). Both groups experienced substantial liberalization, with the OECD average falling from 0.13 to 0.06 and the non-OECD average declining from 0.33 to 0.18.

Panel (b) presents the OECD FDI Regulatory Restrictiveness Index, which measures statutory restrictions on foreign direct investment across countries. Both OECD and non-OECD countries have experienced substantial reductions in FDI barriers since 1997. The average FDI restrictiveness index for OECD countries fell from approximately 0.13 to 0.06, while non-OECD countries saw declines from approximately 0.33 to 0.18. These reductions reflect widespread liberalization of foreign investment regimes, including the removal of equity restrictions, screening requirements, and operational restrictions on foreign-owned firms.

C Productivity and Selection: Additional Evidence

This appendix provides additional evidence on the relationship between firm-level productivity and the key outcomes of interest: multinational status and intangible intensity. These results complement Section 2.5 by presenting full regression specifications and examining the robustness of the productivity gradient across alternative measures and controls.

We estimate the following baseline specification:

$$y_{it} = \alpha + \beta \cdot \text{TFP}_{it} + X_{it}'\gamma + \mu_{n(i)} + \delta_t + \varepsilon_{it} \tag{39}$$

where y_{it} denotes the outcome of interest for firm i in year t. Across Tables 7–8, we consider several outcome variables: an indicator for multinational status, the number of foreign countries in which the firm operates subsidiaries, the foreign sales share, and intangible intensity. The regressor TFP_{it} denotes firm-level total factor productivity, estimated using the control function approach of Ackerberg et al. (2015), and residualized by industry-year fixed effects to remove common productivity shocks. The vector X_{it} includes time-varying firm-level controls (e.g., firm size), and $\mu_{n(i)}$ and δ_t denote 2-digit NAICS industry and year fixed effects, respectively. Standard errors are clustered at the industry and year level to account for correlated shocks within industries and across firms over time.

The coefficient of interest, β, captures the productivity gradient for each outcome. A positive estimate indicates that higher-productivity firms exhibit greater multinational activity or higher intangible intensity consistent with the selection mechanisms emphasized in Section 3. In select specifications examining intangible intensity (Table 8), we augment equation (39) with a quadratic term in TFP_{it} to test for nonlinearities in the productivity–intangible intensity relationship.

C.1 Productivity and Multinational Activity

Table 7 examines how firm-level TFP relates to multinational activity along both extensive and intensive margins. Column (1) shows that a one-standard-deviation increase in TFP is associated with a 15 percentage point increase in the probability of being an MNE, controlling only for year fixed effects. This estimate remains economically large (11 percentage points) after adding industry fixed effects and controlling for firm size in column (2), indicating that the productivity–MNE relationship operates within industries and is not driven solely by scale.

Columns (3)–(4) examine the intensive margin using the number of foreign countries in which

firms operate subsidiaries. The coefficient of 4.2 in the preferred specification (column 4) implies that high-productivity firms maintain substantially broader geographic footprints conditional on MNE status. Columns (5)–(6) confirm this pattern using foreign sales shares: more productive firms derive a larger fraction of their revenue from foreign operations.

Table 7: Productivity, Multinational Status, and Global Scope

Dependent Var	MNE Indicator		Foreign Countries		Foreign Sales Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Revenue TFP	0.15***	0.11***	5.40***	4.20***	0.08***	0.06***
	(0.01)	(0.02)	(0.96)	(0.96)	(0.01)	(0.02)
Industry FE	N	Y	N	Y	N	Y
Control for log assets	N	Y	N	Y	N	Y
Observations	24,688	24,688	24,688	24,688	24,688	24,688

Notes: This table reports OLS estimates of the relationship between firm-level total factor productivity (TFP) and measures of multinational activity. The dependent variable in columns (1)–(2) is an indicator equal to one if the firm reports at least one foreign subsidiary in Exhibit 21 of its SEC 10-K filing. The dependent variable in columns (3)–(4) is the number of distinct foreign countries in which the firm operates subsidiaries. The dependent variable in columns (5)–(6) is the share of foreign segment revenue in total consolidated revenue. TFP is estimated using the Ackerberg et al. (2015) control function approach, with cost of goods sold as the proxy variable, and residualized by industry-year fixed effects. All specifications include year fixed effects. Columns (2), (4), and (6) additionally include 2-digit NAICS industry fixed effects and control for log total assets. Standard errors, reported in parentheses, are clustered at the industry and year level. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

These results are consistent with the theoretical prediction that only the most productive firms find it profitable to incur the fixed costs of multinational expansion (Helpman et al., 2004). However, the strength of the productivity gradient, particularly the finding that productivity explains variation in multinational scope even after controlling for firm size, suggests that productivity operates through mechanisms beyond simple scale effects. In our model, this reflects the complementarity between productivity and intangible capital: high-productivity firms have stronger incentives to invest in scalable intangibles, which in turn raises the returns to geographic expansion.

C.2 Productivity and Intangible Intensity

Table 8 documents a strong positive relationship between TFP and intangible intensity. Column (1) shows that a one-standard-deviation increase in TFP is associated with a 6 percentage point increase in intangible intensity. This relationship is robust to the inclusion of industry fixed effects (column 2), indicating that the productivity–intangible relationship operates within rather than across industries.

Column (3) tests for nonlinearities by including both linear and quadratic terms in TFP. The positive and statistically significant coefficient on the quadratic term (0.005) indicates that

Table 8: Productivity and Intangible Intensity

Dependent Var	Intangible Intensity			
	(1)	(2)	(3)	(4)
Revenue TFP	0.06***	0.06***		0.13***
	(0.006)	(0.006)		(0.02)
Revenue TFP (linear)			0.06***	
			(0.006)	
Revenue TFP (quadratic)			0.005**	
			(0.002)	
Industry FE	N	Y	Y	Y
Control for log assets	N	N	N	Y
Observations	24,688	24,688	24,688	24,688

Notes: This table reports OLS estimates of the relationship between firm-level total factor productivity (TFP) and intangible intensity. The dependent variable in all columns is intangible intensity, defined as the ratio of intangible assets (including capitalized R&D following Peters and Taylor 2017) to total assets. TFP is estimated using the Ackerberg et al. (2015) control function approach, with cost of goods sold as the proxy variable, and residualized by industry-year fixed effects. Column (3) includes both linear and quadratic terms in TFP to test for nonlinearities in the productivity–intangible intensity relationship. All specifications include year fixed effects. Columns (2)–(4) additionally include 2-digit NAICS industry fixed effects. Column (4) controls for log total assets. Standard errors, reported in parentheses, are clustered at the industry and year level. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

the productivity–intangible intensity relationship is convex: intangible intensity increases more steeply at higher productivity levels. This pattern is consistent with our theoretical result (Proposition 3) that cost reductions have larger effects for high-productivity firms, implying that the induced increase in intangible intensity should be concentrated among the most productive firms.

Column (4) adds a control for log total assets. The coefficient on TFP increases substantially (from 0.06 to 0.13), indicating that conditioning on firm size strengthens the productivity–intangible intensity relationship. This pattern suggests that, within size categories, more productive firms are substantially more intangible-intensive, consistent with productivity and intangible investment being complements in production.

D Aggregation of TFP Levels and Growth with Intangible Capital

We want to understand how multinational production, the non-rivalry of intangible capital, and related measurement issues affect the level and growth of aggregate TFP. In particular, we are interested in how these forces affect the TFP level and in how aggregate TFP constructed by a bottom-up approach, which takes a properly weighted average of plant-level TFP, compares with

the aggregate-residual approach undertaken by the BEA and BLS. We begin with plant-level Cobb–Douglas technology because this functional form allows us to align the bottom-up measure of aggregate TFP with the aggregate-residual approach in levels. The second subsection then turns to CES growth and the Tornqvist implementation.

D.1 Cobb–Douglas at the Plant Level: The Level of Aggregate TFP

Consider country c. Plants operating in c belong to three groups: domestic non-MNEs, indexed by $\mathcal{I}_{NM,c}$; domestic MNEs headquartered in c, indexed by $\mathcal{I}_{DMNE,c}$; and foreign MNEs that produce in c but choose their intangible capital abroad, indexed by $\mathcal{I}_{FMNE,c}$. Let $\mathcal{I}_c \equiv \mathcal{I}_{NM,c} \cup \mathcal{I}_{DMNE,c} \cup \mathcal{I}_{FMNE,c}$. We use Cobb–Douglas here because it allows the bottom-up and aggregate-residual approaches to align exactly when they are built on the same input concept. For each plant $i \in \mathcal{I}_c$,

$$y_{ic} = a_{ic} k_{ic}^{\alpha} l_{ic}^{\beta} z_i^{\gamma}, \qquad \eta \equiv \alpha + \beta + \gamma < 1, \tag{40}$$

where z_i is the firm-level non-rival intangible stock deployed at plant i. For $i \in \mathcal{I}_{NM,c} \cup \mathcal{I}_{DMNE,c}$, the choice of z_i is made in country c; for $i \in \mathcal{I}_{FMNE,c}$, it is made in the foreign headquarters country. Booked and used intangibles need not coincide:

$$z_{ic}^{booked} = \mu_{ic} z_i, \tag{41}$$

with $\mu_{ic} = 1$ under perfect measurement.

Bottom-up TFP and the aggregate residual. Write

$$Y_c = Y_{NM,c} + Y_{DMNE,c} + Y_{FMNE,c} = \sum_{i \in \mathcal{I}_c} y_{ic},$$

$$K_c = \sum_{i \in \mathcal{I}_c} k_{ic}, \qquad L_c = \sum_{i \in \mathcal{I}_c} l_{ic}, \qquad Z_c^{used} = \sum_{i \in \mathcal{I}_c} z_i, \qquad Z_c^{booked} = \sum_{i \in \mathcal{I}_c} z_{ic}^{booked},$$

where, for instance, $Y_{FMNE,c} \equiv \sum_{i \in \mathcal{I}_{FMNE,c}} y_{ic}$. Let $\omega_{k,ic} \equiv k_{ic}/K_c$, $\omega_{l,ic} \equiv l_{ic}/L_c$, and $\omega_{z,ic} \equiv z_i/Z_c^{used}$. Substituting $k_{ic} = \omega_{k,ic} K_c$, $l_{ic} = \omega_{l,ic} L_c$, and $z_i = \omega_{z,ic} Z_c^{used}$ into (40) and summing over the three groups yields

$$Y_c = \left(\sum_{i \in \mathcal{I}_c} a_{ic} \omega_{k,ic}^{\alpha} \omega_{l,ic}^{\beta} \omega_{z,ic}^{\gamma} \right) K_c^{\alpha} L_c^{\beta} (Z_c^{used})^{\gamma}. \tag{42}$$

Hence the bottom-up TFP level is

$$A_c^{BU} \equiv \sum_{i \in \mathcal{I}_c} a_{ic} \omega_{k,ic}^{\alpha} \omega_{l,ic}^{\beta} \omega_{z,ic}^{\gamma}. \tag{43}$$

If the aggregate residual is constructed with the same production-relevant intangible aggregate,

$$A_c^{RES,used} \equiv \frac{Y_c}{K_c^{\alpha} L_c^{\beta} (Z_c^{used})^{\gamma}}, \tag{44}$$

then (42) implies

$$A_c^{BU} = A_c^{RES,used}. \tag{45}$$

Equation (45) is the natural starting point. The bottom-up object and the aggregate residual align exactly as long as they use the same production boundary. The interesting questions are therefore not about algebra, but about which intangible input and which income concept enter the aggregate statistic.

The productivity implication of multinational scope, non-rivalry, the effective user cost of z. The theory section implies that non-rivalry lowers the shadow cost of intangibles at the plant level when a firm operates in more locations. In the baseline theory, if a firm serves J_i locations, its first-order condition for z_i is

$$J_i paF_z(z_i, l_i) = MC_c.$$

Equivalently, each plant behaves as if it faced the user cost

$$q_i^z(J_i) \equiv \frac{MC_c}{J_i}. \tag{46}$$

Equivalently,

$$q_i^z(J_i) \equiv (1 + \tau_z(J_i)) MC_c, \qquad 1 + \tau_z(J_i) = \frac{1}{J_i}, \qquad \tau_z'(J_i) < 0. \tag{47}$$

For a domestic non-MNE, $J_i = 1$ and $q_i^z(J_i) = MC_c$. For a domestic MNE, $J_i > 1$ and the intangible decision is made in c. For a foreign MNE plant in c, the same formula describes the shadow price of intangible services used locally, but that shadow price is chosen abroad. This is the sense in which multinational scope and ownership jointly shape within-country production.

With competitive factor markets for rival inputs, the plant-level conditions are

$$k_{ic} = \frac{\alpha}{R_c} y_{ic}, \qquad l_{ic} = \frac{\beta}{W_c} y_{ic}, \qquad z_i = \frac{\gamma}{q_i^z(J_i)} y_{ic}. \tag{48}$$

These three relationships hold for domestic non-MNEs and for domestic plants of domestic MNEs. For domestic plants of foreign MNEs, intangible capital is determined in the foreign headquarters country, denoted c'. Under symmetric countries, the same choice satisfies

$$z_i = \frac{\gamma}{q_i^z(J_i)} y_{ic} = \frac{\gamma}{q_i^z(J_i)} y_{ic'},$$

where c' is the headquarters country of firm i. Substituting (48) into (40) yields

$$y_{ic} = \Psi_c a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}} = \Psi_c a_{ic}^{\frac{1}{1-\eta}} J_i^{\frac{\gamma}{1-\eta}}, \tag{49}$$

where

$$\Psi_c \equiv \left(\frac{\alpha^\alpha \beta^\beta \gamma^\gamma}{R_c^\alpha W_c^\beta MC_c^\gamma} \right)^{\frac{1}{1-\eta}} > 0.$$

The same firm-level non-rival stock therefore has a lower effective user cost and a larger plant-level revenue product when it can be deployed over more affiliates. This scope dependence is also the source of real misallocation in the aggregate TFP level. Equation (49) implies

$$Y_c = \Psi_c \sum_{i \in \mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}},$$

$$K_c = \frac{\alpha}{R_c} Y_c, \qquad L_c = \frac{\beta}{W_c} Y_c, \qquad Z_c^{used} = \frac{\gamma}{MC_c} \Psi_c \sum_{i \in \mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}-1}.$$

Hence

$$\omega_{k,ic} = \omega_{l,ic} = \frac{y_{ic}}{Y_c} = \frac{a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}}}{\sum_{n \in \mathcal{I}_c} a_{nc}^{\frac{1}{1-\eta}} \left(\frac{q_n^z(J_n)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}}},$$

$$\omega_{z,ic} = \frac{a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}-1}}{\sum_{n \in \mathcal{I}_c} a_{nc}^{\frac{1}{1-\eta}} \left(\frac{q_n^z(J_n)}{MC_c} \right)^{-\frac{\gamma}{1-\eta}-1}}.$$

Substituting these shares into (43) gives

$$A_c^{BU} = \frac{\left(\sum_{i\in\mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c}\right)^{-\frac{\gamma}{1-\eta}}\right)^{1-\eta+\gamma}}{\left(\sum_{i\in\mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c}\right)^{-\frac{\gamma}{1-\eta}-1}\right)^{\gamma}}. \tag{50}$$

This object already incorporates the three-way composition of production in country c. Domestic non-MNEs enter with $q_i^z(J_i) = MC_c$. Domestic MNEs enter with lower $q_i^z(J_i)$ when they serve more foreign locations. Foreign MNE plants also enter with low effective user costs, but those costs reflect intangible decisions taken abroad. What matters for production is not the ownership of the firm per se, but the dispersion in effective user costs across plants operating in c.

The frictionless benchmark is the case in which $q_i^z(J_i) = q^z$ is common across all plants in country c. Then the common factor cancels from (50), so

$$A_c^{FB} = \left(\sum_{i\in\mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}}\right)^{1-\eta}. \tag{51}$$

Hence only dispersion in effective user costs matters for real misallocation:

$$\frac{A_c^{BU}}{A_c^{FB}} = \frac{\left(\sum_{i\in\mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c}\right)^{-\frac{\gamma}{1-\eta}}\right)^{1-\eta+\gamma}}{\left(\sum_{i\in\mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}} \left(\frac{q_i^z(J_i)}{MC_c}\right)^{-\frac{\gamma}{1-\eta}-1}\right)^{\gamma} \left(\sum_{i\in\mathcal{I}_c} a_{ic}^{\frac{1}{1-\eta}}\right)^{1-\eta}}. \tag{52}$$

The key point is that non-rivalry does not merely raise average productivity. By lowering $q_i^z(J_i)$ differentially across non-MNEs, domestic MNEs, and foreign MNE plants, it changes the allocation of output and inputs within the country.

Accounting mismeasurement: booked versus used intangibles. The production-consistent residual in (44) uses $Z_c^{used} = \sum_{i\in\mathcal{I}_c} z_i$. National accounts instead observe a booked stock,

$$Z_c^{booked} = \sum_{i\in\mathcal{I}_c} \mu_{ic} z_i = \sum_{i\in\mathcal{I}_{NM,c}} \mu_{ic} z_i + \sum_{i\in\mathcal{I}_{DMNE,c}} \mu_{ic} z_i + \sum_{i\in\mathcal{I}_{FMNE,c}} \mu_{ic} z_i.$$

For foreign MNE plants in c, μ_{ic} is naturally low because the intangible asset is typically booked in the headquarters country even though its services are used in production in c. The residual based

on booked intangibles is

$$A_c^{RES,booked} \equiv \frac{Y_c}{K_c^{\alpha} L_c^{\beta} (Z_c^{booked})^{\gamma}}. \tag{53}$$

Using (42),

$$A_c^{RES,booked} = A_c^{BU} \left(\frac{Z_c^{used}}{Z_c^{booked}} \right)^{\gamma}. \tag{54}$$

Equivalently,

$$\frac{A_c^{RES,booked}}{A_c^{BU}} = \left(\frac{Z_c^{used}}{Z_c^{booked}} \right)^{\gamma}. \tag{55}$$

This wedge is purely an accounting wedge. The bottom-up object is unchanged because it aggregates the input actually used in production. The booked-stock residual differs because it replaces the used stock with an ownership- and reporting-based measure. With inward multinational production, one typically has $Z_c^{used} > Z_c^{booked}$, so the booked-capital residual overstates production-based TFP.

GNP TFP through the aggregate-residual approach. The previous objects are production-based and therefore naturally tied to GDP. When intangible capital is non-rival, however, a stock created by a domestic firm generates income in multiple countries at once. A TFP measure based on GNP rather than GDP therefore better captures the nature of non-rivalry, because it follows the income generated by the domestically owned intangible stock rather than assigning all returns according to the location of production. To move to GNP, one must switch from a geographic production boundary to an ownership-based income boundary. Let

$$Y_{DMNE,-c} \equiv \sum_{i \in \mathcal{I}_{DMNE,c}} \sum_{m \neq c} y_{im}$$

denote output produced abroad by domestic MNEs headquartered in c. Then the GNP income of country c can be written exactly as

$$\begin{aligned} GNP_c &= Y_{NM,c} + Y_{DMNE,c} + \sum_{i \in \mathcal{I}_{FMNE,c}} (y_{ic} - q_i^z(J_i) z_i) + \sum_{i \in \mathcal{I}_{DMNE,c}} \sum_{m \neq c} q_i^z(J_i) z_i \\ &= Y_c - \sum_{i \in \mathcal{I}_{FMNE,c}} q_i^z(J_i) z_i + \sum_{i \in \mathcal{I}_{DMNE,c}} \sum_{m \neq c} q_i^z(J_i) z_i. \end{aligned} \tag{56}$$

The first term is the value added of domestic non-MNEs and domestic plants of domestic MNEs. The second subtracts the income generated by foreign-owned intangibles from foreign MNE plants operating in c. The third adds the intangible income that domestic MNEs earn abroad on the same

non-rival stock owned in c.

Using the plant-level condition in (48), $q_i^z(J_i)z_i = \gamma y_{ic}$ at every plant, so (56) becomes

$$GNP_c = Y_c - \gamma Y_{FMNE,c} + \gamma Y_{DMNE,-c}. \tag{57}$$

This ownership adjustment cannot be recovered from the bottom-up TFP measure alone, because the bottom-up object is defined over plants producing in country c. Once the target is GNP rather than GDP, the relevant aggregate input concept must also switch from used intangibles in country c to the booked stock owned by residents of c:

$$Z_c^{own} \equiv \sum_{i \in \mathcal{I}_{NM,c} \cup \mathcal{I}_{DMNE,c}} z_{ic}^{booked}. \tag{58}$$

The relevant aggregate residual is therefore

$$A_c^{RES,GNP} \equiv \frac{GNP_c}{K_c^{\alpha} L_c^{\beta} (Z_c^{own})^{\gamma}}. \tag{59}$$

Combining (44), (58), and (57) yields

$$A_c^{RES,GNP} = A_c^{BU} \left(1 - \gamma \frac{Y_{FMNE,c}}{Y_c} + \gamma \frac{Y_{DMNE,-c}}{Y_c}\right) \left(\frac{Z_c^{used}}{Z_c^{own}}\right)^{\gamma}. \tag{60}$$

Equation (60) summarizes the logic of the subsection. A_c^{BU} is the production-based, bottom-up TFP level. $A_c^{RES,used}$ aligns with it exactly. $A_c^{RES,GNP}$ differs for two reasons: it applies an ownership correction to the numerator by subtracting foreign-owned intangible income earned at home and adding domestically owned intangible income earned abroad, and it uses the booked intangible stock owned by residents rather than the non-rival stock used in production in country c.

D.2 Bottom-Up and Aggregate-Residual TFP Growth under CES

When production combines non-rival intangible capital with rival local inputs, the central question is whether the aggregate residual recovers the output-share-weighted growth of plant technology. Under CES production, that equivalence is not automatic. The gap is economically relevant for two reasons. First, multinational expansion allows the same stock of intangible capital to be deployed across more plants, which changes plant-level elasticities when $\sigma > 1$. Second, national accounts may measure booked rather than used intangible capital, so aggregate residuals can absorb changes in accounting location rather than technology.

Consider plant i:

$$y_i = A_i \left[\mu z_i^{\frac{\sigma-1}{\sigma}} + (1-\mu) r_i^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\theta\sigma}{\sigma-1}}, \qquad r_i \equiv k_i^{\alpha} l_i^{1-\alpha}, \tag{61}$$

where z_i is non-rival intangible capital, r_i is the rival input composite, $\mu \in (0,1)$ is the CES weight on z_i, and $\sigma > 1$ is the elasticity of substitution between the non-rival and rival inputs. Let $m \in \{k, l, z\}$ index inputs and define plant-level output elasticities by

$$\varepsilon_{im} \equiv \frac{\partial \ln y_i}{\partial \ln x_{im}}.$$

Under cost minimization, these elasticities coincide with plant-level cost shares.

Plant-level Divisia identity. Totally differentiating (61) gives

$$d \ln y_i = d \ln A_i + \sum_m \varepsilon_{im} \, d \ln x_{im}, \tag{62}$$

and therefore

$$d \ln A_i = d \ln y_i - \sum_m \varepsilon_{im} \, d \ln x_{im}. \tag{63}$$

This is the plant-level technology-growth object that a bottom-up measure should aggregate.

Bottom-up and aggregate-residual growth. With homogeneous output, $Y = \sum_i y_i$ and

$$d \ln Y = \sum_i s_i \, d \ln y_i, \qquad s_i \equiv \frac{y_i}{Y}. \tag{64}$$

Define bottom-up aggregate TFP growth as

$$d \ln A^{BU} \equiv \sum_i s_i \, d \ln A_i. \tag{65}$$

Substituting (62) into (64) yields

$$d \ln Y = d \ln A^{BU} + \sum_m \sum_i s_i \varepsilon_{im} \, d \ln x_{im}. \tag{66}$$

By contrast, the aggregate residual is

$$d\ln A^{RES} \equiv d\ln Y - \sum_m \bar{s}_m \, d\ln X_m, \qquad d\ln X_m = \sum_i \lambda_{im} \, d\ln x_{im}, \qquad \lambda_{im} \equiv \frac{x_{im}}{X_m}, \tag{67}$$

where $\bar{s}_m$ denotes the aggregate income share of input m.

For rival inputs, the aggregate quantity is unambiguous. For non-rival intangibles, the relevant distinction is between intangible capital used in production and intangible capital booked in the national accounts. Let

$$Z_t^{booked} = \overline{\mu}_t Z_t^{used},$$

where $\overline{\mu}_t$ is the aggregate booked-to-used ratio. Then

$$d\ln Z_t^{booked} = d\ln Z_t^{used} + d\ln \overline{\mu}_t. \tag{68}$$

When multinational production expands, the same stock of intangible capital is deployed across more foreign plants, while the booked location of that capital need not move with production. Changes in $\overline{\mu}_t$ therefore capture accounting reallocation rather than technological change.

Main result: the wedge between the two growth measures. Combining (66) and (67) gives

$$d\ln A^{RES} - d\ln A^{BU} = \sum_m \sum_i \left(s_i \varepsilon_{im} - \bar{s}_m \lambda_{im}\right) d\ln x_{im} - \bar{s}_z \, d\ln \overline{\mu}. \tag{69}$$

Equation (69) shows that the aggregate residual differs from bottom-up technology growth for four conceptually distinct reasons: plant-level elasticities may be heterogeneous; aggregate income shares may fail to equal the relevant average elasticities; input growth may be concentrated in firms whose output weights differ from their input weights; and booked intangible capital may differ from the intangible capital actually used in production.

The first and third channels are particularly relevant in an economy with multinational expansion. Because intangible capital is non-rival, firms that expand across plants deploy the same z more broadly. When $\sigma > 1$, that changes the relative importance of z and raises the output elasticity of intangibles at plants with high intangible intensity. If the firms expanding intangible use are also the firms with high output weights, bottom-up growth assigns more of aggregate output growth to technology than does an aggregate residual based on common shares. The final term in (69) is different in nature: it is a pure accounting wedge. If the booked-to-used ratio falls, so that $d\ln \overline{\mu} < 0$, the aggregate residual mechanically overstates TFP growth because measured

intangible input growth understates the growth of intangible services used in production.

To separate the economic channels, rewrite the contribution of input m as

$$\mathcal{W}_m = \underbrace{\sum_i s_i(\varepsilon_{im} - \bar{\varepsilon}_m) d\ln x_{im}}_{\text{elasticity heterogeneity}} + \underbrace{(\bar{\varepsilon}_m - \bar{s}_m) d\ln X_m}_{\text{share-elasticity wedge}} + \underbrace{\bar{\varepsilon}_m \sum_i (s_i - \lambda_{im}) d\ln x_{im}}_{\text{composition/reallocation}}, \tag{70}$$

where $\bar{\varepsilon}_m \equiv \sum_i s_i \varepsilon_{im}$. Then

$$d\ln A^{RES} - d\ln A^{BU} = \sum_m \mathcal{W}_m - \bar{s}_z \, d\ln \overline{\mu}.$$

The sign of the first term is positive when firms with high elasticities also expand input m rapidly. The sign of the third term is positive when input growth is concentrated in firms whose output shares exceed their shares in the aggregate stock of that input. Equality between the two growth measures therefore requires more than common elasticities: it requires no elasticity heterogeneity, no share-elasticity mismatch, no composition wedge, and a constant booked-to-used ratio.

Tornqvist: key takeaway. In discrete time, the Tornqvist index is designed to approximate the continuous-time Divisia residual with observable data: it uses midpoint income shares to remove the contribution of input growth from output growth and thereby recover aggregate TFP growth. It therefore tries to reproduce, with aggregate quantities and shares, the technology-growth object that the bottom-up approach constructs directly from plant-level TFP. The gap between Tornqvist-based TFP growth and bottom-up TFP growth is small when share movements mainly reflect smooth CES substitution. That gap widens when income shares move because multinational expansion reallocates non-rival intangible capital across plants, or because booked intangible capital diverges from intangible capital used in production. In those cases, the Tornqvist index inherits the same non-technology wedges as the aggregate-residual approach.

Discrete-time derivation. In discrete time, Divisia is implemented by Tornqvist:

$$\Delta \ln A_t^{RES} \approx \Delta \ln Y_t - \sum_m \frac{1}{2} \left(s_{m,t} + s_{m,t-1} \right) \Delta \ln X_{m,t}. \tag{71}$$

This is the logic used by BLS-style aggregate TFP measurement: chain-type real output growth net of a Tornqvist index of combined inputs with income-share weights.

Let midpoint objects be

$$\bar{s}_{i,t} \equiv \frac{s_{i,t} + s_{i,t-1}}{2}, \qquad \bar{\lambda}_{im,t} \equiv \frac{\lambda_{im,t} + \lambda_{im,t-1}}{2}, \qquad \bar{\varepsilon}_{im,t} \equiv \frac{\varepsilon_{im,t} + \varepsilon_{im,t-1}}{2},$$

and define

$$\Delta \ln A_t^{TOR} \equiv \Delta \ln Y_t - \sum_m \bar{s}_{m,t} \Delta \ln X_{m,t}.$$

Then a discrete-time analogue of (69) is

$$\Delta \ln A_t^{TOR} - \Delta \ln A_t^{BU} = \sum_m \sum_i \left(\bar{s}_{i,t} \bar{\varepsilon}_{im,t} - \bar{s}_{m,t} \bar{\lambda}_{im,t} \right) \Delta \ln x_{im,t} + \mathcal{R}_t, \tag{72}$$

where $\mathcal{R}_t$ is a second-order discretization remainder. To derive it, define

$$g_{im}(u) \equiv s_i(u) \varepsilon_{im}(u) - s_m(u) \lambda_{im}(u), \qquad u \in [t-1, t].$$

Integrating the continuous-time wedge over the interval $[t-1, t]$ gives

$$\Delta \ln A_t^{RES} - \Delta \ln A_t^{BU} = \sum_m \sum_i \int_{t-1}^{t} g_{im}(u) \, d \ln x_{im}(u).$$

The Tornqvist approximation replaces the path of $g_{im}(u)$ by its midpoint proxy

$$\bar{g}_{im,t} \equiv \bar{s}_{i,t} \bar{\varepsilon}_{im,t} - \bar{s}_{m,t} \bar{\lambda}_{im,t},$$

so that

$$\sum_m \sum_i \int_{t-1}^{t} g_{im}(u) \, d \ln x_{im}(u) = \sum_m \sum_i \bar{g}_{im,t} \Delta \ln x_{im,t} + \sum_m \sum_i \int_{t-1}^{t} \left[g_{im}(u) - \bar{g}_{im,t} \right] d \ln x_{im}(u).$$

Hence

$$\mathcal{R}_t \equiv \sum_m \sum_i \int_{t-1}^{t} \left[g_{im}(u) - \bar{g}_{im,t} \right] d \ln x_{im}(u). \tag{73}$$

If shares, elasticities, and input growth vary smoothly over the interval, then $g_{im}(u) - \bar{g}_{im,t}$ is first order in the length of the interval, while $d \ln x_{im}(u)$ is also first order, implying that $\mathcal{R}_t$ is second

order. Rearranging as above gives

$$\Delta \ln A_t^{TOR} - \Delta \ln A_t^{BU} := \sum_m \Bigg[\sum_i \bar{s}_{i,t}(\bar{\varepsilon}_{im,t} - \bar{\varepsilon}_{m,t}) \Delta \ln x_{im,t} + (\bar{\varepsilon}_{m,t} - \bar{s}_{m,t}) \Delta \ln X_{m,t} + \bar{\varepsilon}_{m,t} \sum_i (\bar{s}_{i,t} - \bar{\lambda}_{im,t}) \Delta \ln x_{im,t} \Bigg] + \mathcal{R}_t, \tag{74}$$

with $\bar{\varepsilon}_{m,t} \equiv \sum_i \bar{s}_{i,t} \bar{\varepsilon}_{im,t}$. If time-varying shares mainly reflect CES substitution, the first-order wedge terms are small and Tornqvist differs from bottom-up growth only through the second-order remainder. If instead share movements reflect wedges, booking shifts, or mismeasured user costs, then Tornqvist loads those forces directly into measured aggregate TFP growth.