Quality and Outsourcing in Advanced Manufacturing Networks *

Jen-Kuan Wang
Penn State

Jin-Tan Liu
National Taiwan University
NBER

November 8, 2026

[Link to the latest version]

Abstract

In production networks, upstream markups give downstream firms an incentive to make inputs in-house, whereas a comparative advantage in their core production pushes them to outsource. We formalize this make-or-buy trade-off in a bipartite network model with input-output quality complementarities, endogenous entry, and endogenous quality choices, both upstream and downstream. Fit to a large set of moments describing Taiwanese advanced manufacturers, the model explains two patterns we document in firm-to-firm transaction data: high-quality firms outsource more, and they buy disproportionately from high-quality suppliers. We then use the estimated model for counterfactual policy experiments. Policies limiting the market power of upstream firms can reduce welfare despite lessening the distortion from double marginalization. The reason is that markups also finance suppliers' quality investment; removing them causes upstream firms to exit, forcing downstream firms into inefficient in-house production and lowering output quality. Tax credits for input purchases can instead raise welfare, because they address the distortion from upstream markups without shifting input production toward less efficient firms.

JEL Codes: F12, L14, L25, O30

Keywords: vertical specialization, quality upgrading, markups, industrial policy

*Wang (corresponding author): Department of Economics, Pennsylvania State University, 303 Kern Building, University Park, PA 16802, USA (jkwang@psu.edu). Liu: National Taiwan University and NBER (liujt@ntu.edu.tw). This research was conducted in collaboration with the Fiscal Information Agency, Ministry of Finance, Taiwan. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Fiscal Information Agency, Ministry of Finance, or any institution with which the authors are affiliated. This paper is Wang's job market paper; he is grateful for the mentorship of James Tybout, Mark Roberts, and Stephen Yeaple, and for feedback from Maria-Jose Carreras-Valle, Marc Henry, Seula Kim, Kala Krishna, Eric San Miguel, Conor Ryan, Kai-Jie Wu, Ruilin Zhou, Toshiaki Komatsu, Joonkyo Hong, Oliver Kim, and participants at the Dissertation Fellow Workshop at the Federal Reserve Bank of St. Louis. All remaining errors are our own.

1 Introduction

Economists have long recognized that the distortionary effects of monopolistic pricing cascade down vertical supply chains.[1] Since Spengler (1950), this double marginalization has been cited as an argument for the in-house production of intermediate inputs. But in-house production may not be advisable if it means giving up the comparative advantage and scale of specialized upstream suppliers. This consideration is particularly acute in advanced manufacturing, where output quality hinges on input quality and upgrading is costly. The gains from outsourcing input production can then outweigh the cost of the markups suppliers charge. Understanding this make-or-buy trade-off in quality-critical supply chains is essential for designing policies aimed at addressing market imperfections while fostering quality investment.

We develop a bipartite network model that embodies these benefits and costs of in-house production, fitting it to detailed micro data on Taiwanese advanced manufacturers. Two features of our model are key. First, there are quality complementarities between inputs and outputs, as in Fieler, Eslava, and Xu (2018) and Demir et al. (2024). That is, better products require better inputs.[2] Second, downstream firms are heterogeneous in two dimensions: their ability to produce high-quality output and their ability to produce high-quality intermediate inputs. The two need not go together: a firm that excels at its final product is often less efficient at making high-quality inputs. Together, these assumptions imply that firms' incentives to avoid double marginalization by producing inputs in-house are sometimes dominated by their incentives to take advantage of the greater productive efficiency of upstream firms, and that this happens relatively frequently among high-quality producers. The model thus helps explain the vertical specialization seen in advanced manufacturing today, for example, why Nvidia outsources the production of its most advanced chips to Taiwan Semiconductor Manufacturing Company (TSMC).

To study the relationship between firms' quality choices and sourcing decisions, we exploit several Taiwanese administrative datasets. Specifically, we merge value-added tax (VAT) records covering the universe of domestic firm-to-firm transactions with corporate tax filings and customs records for 2013–2022. The VAT records allow us to measure how much of its production chain a firm outsources, following Boehm and Oberfield (2020, 2023). The tax filings and customs records allow us to construct two measures of output quality, based on R&D spending (Sutton 2001; Kugler and Verhoogen 2012) and on the demand-residual quality measure of Khandelwal (2010), respectively.

We focus on four advanced manufacturing sectors that jointly account for 80% of Taiwan's

1. The observation first appears in Cournot (1838).

2. Formally, complementarity means that higher output quality raises the marginal return to input quality. See Costinot (2009) for a detailed discussion.

manufacturing R&D. Their combination of intensive quality investment and fragmented supply chains makes them a useful setting for studying the trade-off of interest. Firms in these sectors also source most of their inputs domestically, allowing us to observe their supplier relationships in the transaction data. At the same time, these sectors are central to modern industrial policy. A 2023 amendment to Taiwan's Statute for Industrial Innovation, for example, provides a 5% tax credit for expenditures on advanced-process equipment.[3]

We estimate the model by the method of simulated moments, matching the joint distribution of firms' sourcing choices, quality choices, and vertical specialization. Two patterns in the data are key for identification: more vertically specialized firms produce higher-quality output, and more specialized firms buy from high-quality suppliers. These pin down the correlation between firms' abilities to produce input and output quality, and the parameter that governs the strength of the complementarity between inputs and outputs, respectively. Intuitively, if the correlation between the two abilities were too strongly positive, high-quality firms would outsource relatively less, as their larger scale allows them to afford in-house quality investment. And without sufficiently strong input-output quality complementarity, high-quality firms would outsource only slightly more and would no longer sort toward higher-quality suppliers.

Upstream markups distort the behavior of downstream firms in several ways. Relative to the social planner's solution, they induce these firms to produce more of their inputs than is optimal, and they cause them to adjust their input and output quality choices downward. But upstream markups are not unequivocally bad. They also reward suppliers for investing in quality, since suppliers capture a portion of the associated rents (Schumpeter 1934; Grossman and Hart 1986). We close our analysis by performing counterfactual experiments that demonstrate how these margins of adjustment respond to two kinds of policies: those that reduce the market power of upstream firms, and those that subsidize purchases of intermediate inputs.

The first policy proves to be disastrous. Squeezing upstream margins induces exit among upstream firms, forcing downstream firms to produce the inputs they need in-house, often inefficiently. In the limit, when upstream firms are forced to price at marginal cost, they all exit, output quality declines, and real consumption falls by 40%. The second policy is far more effective. A credit on input purchases lowers downstream buyers' prices while preserving upstream suppliers' operating profits. Applied at Taiwan's 5% rate for advanced-process equipment and financed by a lump-sum charge on households, it raises real consumption by 2.7%.[4]

Taken together, our findings show that quality upgrading and vertical specialization can

3. The amendment added Article 10-2 to the Statute, often called Taiwan's Chips Act. Similar laws elsewhere subsidize investment in advanced manufacturing facilities and equipment, including the U.S. CHIPS and Science Act (2022), the European Chips Act (2023), and Korea's K-Chips Act (2023). The text of the Statute, in Chinese, is available from the Ministry of Justice's Laws and Regulations Database.

4. A 40% credit, which offsets the markup in full, raises real consumption by 18%.

reinforce one another. Empirically, we provide direct firm-to-firm evidence that more specialized manufacturers produce higher-quality output and source disproportionately from quality-investing suppliers. Our model explains these patterns through comparative advantage amplified by input-output quality complementarity and quantifies the gains and costs of outsourcing. The central lesson is that suppliers' markups are both a distortion of downstream production and a reward for upstream quality investment. Policies that reduce the costs of outsourcing must therefore preserve the incentives for quality investment by upstream suppliers.

Contributions and Related Literature. This paper speaks to three strands of literature: on firm boundaries and quality investment, on buyer-seller sorting in production networks, and on markups and industrial policy. We make three contributions. Empirically, using the universe of domestic firm-to-firm transactions in Taiwan's advanced manufacturing, we provide, to our knowledge, the first direct evidence that firms outsourcing more of their inputs produce higher-quality output, invest more in quality, and source from suppliers that do the same. Our model explains these facts: comparative advantage and complementarity between input and output quality lead high-quality firms to outsource despite suppliers' markups, which in turn fund their suppliers' upgrading. Quantitatively, we measure the misallocation these markups create and show that a credit on input purchases raises welfare, whereas removing the markups lowers it.

Since Coase (1937), the literature on the boundary of the firm has predicted that frictions arising in input markets but not within the firm push production in-house. Suppliers' markups in our setting can stand in for the costs of using the market that result from transaction costs (Riordan and Williamson 1985), contracting frictions and residual rights of control (Grossman and Hart 1986; Antràs 2003; Boehm and Oberfield 2020), coordination costs (Becker and Murphy 1992), and suppliers' market power (Dhyne, Kikkawa, and Magerman 2022; Cajal-Grossi, Macchiavello, and Noguera 2023; Tsai 2025). Rather than take a stance on the source of the markup, we use evidence and a model to show how quality investment shapes vertical specialization in the presence of these costs.

Our findings provide a useful contrast to earlier work that links capital intensity or quality to less outsourcing. Antràs (2003) shows that the share of intrafirm trade rises with the capital intensity of the sector, and in Peruvian fish meal, Hansman et al. (2020) provide causal evidence that a higher premium for output quality leads firms to integrate their suppliers. In Taiwan's advanced manufacturing, we find the opposite: high-quality firms are more vertically specialized. Our model attributes the reversal to input-output quality complementarity on top of comparative advantage: high-quality firms are not necessarily good at making premium inputs, for which they have particularly high demand. Closer to our findings, Acemoglu et al. (2010) show, in theory and in U.K. data, that integration becomes less likely as the sup-

plier's technology intensity rises, especially when its input is a large share of the buyer's costs.

Our emphasis on quality complementarity connects the paper to a growing literature on buyer-seller matching and sorting in international and domestic production networks (Eaton et al. 2025; Eaton, Kortum, and Kramarz 2022; Eaton et al. 2022; Bernard and Moxnes 2018; Bernard et al. 2022; Lim 2018; Huneeus 2018; Miyauchi 2024; Arkolakis, Huneeus, and Miyauchi 2023; Demir et al. 2024). Closest to our work, Demir et al. (2024) bring the O-ring logic of Kremer (1993) to production networks and document positive assortative matching on wages in Turkish data: higher-paying firms source from higher-paying suppliers, consistent with complementarity between input and output quality. Viratyosin (2026) documents buyer-seller sorting on R&D intensity in U.S. data and studies its implications for growth.

Relative to these papers, we emphasize that complementarity can shape not only whom firms buy from, but also how much they outsource. Bijnens, Humlum, and Vanormelingen (2026) show, using Belgian firm-to-firm data, that a firm's investment in capital complements the skilled labor employed by its suppliers. In their setting, complementarity also crosses firm boundaries, linking a firm's capital to its suppliers' skilled labor; in ours, it links the firm's output quality to its suppliers' input quality. We contribute by quantifying how such complementarity amplifies the gains from outsourcing on top of comparative advantage.

In our estimates, suppliers' markups discourage outsourcing and raise costs most for high-quality firms, which rely more heavily on external inputs. We contribute by quantifying the resulting misallocation in advanced manufacturing and evaluating how policy can mitigate it without discouraging suppliers' quality investment. Liu (2019) shows that distortions compound through input-output linkages, providing a rationale for targeting upstream sectors with industrial policy. Our results highlight an additional consideration: policy can address the upstream markup distortion through a credit on input purchases, lowering buyers' costs while preserving the supplier margins that support quality investment.

The key lesson is that suppliers' markups not only create distortions but also reward suppliers for investing in quality. This is consistent with the literature on hold-up problems, in which suppliers underinvest ex ante when they cannot retain enough of the returns ex post (Grossman and Hart 1986). Macchiavello et al. (2026) make a related point for Colombian coffee, where farmers invest in quality when an international buyer requires the exporter to pay them a quality premium. In our setting, complementarity between input and output quality makes these incentives particularly important: suppliers' quality investment also raises the return to their buyers' own quality upgrading. Consequently, eliminating suppliers' markups can undermine quality investment throughout the supply chain, even as it reduces double marginalization.

The rest of the paper is organized as follows. Section 2 introduces the data and the measure of vertical specialization, and Section 3 presents the motivating facts. Section 4 lays out the

theoretical framework, Section 5 details the estimation, and Section 6 presents the quantitative exercises. Section 7 concludes.

2 Data and the Measure of Vertical Specialization

2.1 Data

Our analysis is based on several tax-related datasets administered by Taiwan's Ministry of Finance and its Fiscal Information Agency, spanning 2013–2022. All datasets are linked through a common firm identifier. At the center of our analysis are firm-to-firm transaction data from Taiwan's value-added tax (VAT) system, which records buyer-seller transaction values bimonthly (up to six observations per year). We aggregate these transactions to annual buyer-seller totals. We link these data to firm-level financial statements from corporate income tax filings, which report sales, detailed factor input usage, quality-related expenditures such as R&D and advertising, and intangible asset holdings. We further merge customs records on imports and exports, which report values and quantities (in kilograms) by 6-digit Harmonized System (HS6) product; we aggregate these records to the firm-HS6-country level at an annual frequency. Finally, we supplement these data with personal income tax records, which identify each employee's employer and allow us to construct firm-level payroll and headcounts.

Two institutional features make these data reliable and comprehensive. First, firms are required to record sales using distinct receipt formats for business-to-business (B2B), business-to-consumer, and export transactions, so the three components of sales are observed separately. Following Dhyne et al. (2021), we define total sales as their sum. Second, the firm-to-firm records have no minimum reporting threshold: the transaction value of every buyer-seller pair is reported regardless of size.[5]

Preparing the network data. For the main analysis of this paper, we impose a sequence of cleaning and selection criteria to construct internally consistent production networks (Dhyne et al. 2021). We first retain firms with positive total sales and positive labor costs. After this initial selection, we recompute network sales using the updated sample. Some firms lose all customers as a result and record zero sales; these firms are dropped in subsequent rounds. We apply this procedure iteratively until we reach a network system in which all remaining firms have input and output information that is internally consistent. As shown in Chen et al. (2026), while the cleaned network sample retains only 37% of tax IDs, it covers over 83% of total sales

5. Chen et al. (2026) document the institutional details of the Taiwanese value-added tax system, verify the consistency of business-to-business sales across the transaction-level and tax-return data, and compare the combined dataset against official aggregate statistics.

and costs relative to the raw data, which in turn closely track the aggregate statistics published by Taiwan's statistical bureau. Using the cleaned network data, we construct firm-level supplier and customer characteristics weighted by network transaction share.

Next, we apply the following sample selections prior to analysis. First, to shed light on the relationship between quality investment and the vertical structure of the firm, we restrict attention to firms in the four advanced manufacturing sectors: Electronic Parts and Components; Computers, Electronic and Optical Products; Electrical Equipment; and Machinery and Equipment. These industries have the largest potential for quality variation as they account for around 80% of R&D spending in Taiwan's manufacturing sector during the data period. At the same time, their domestic supply chains are fragmented into many specialized firms.[6] Second, to minimize concerns about misreporting in tax data, we retain only privately owned firms subject to regular audits by the tax authority and with annual sales above the standard value-added tax eligibility threshold.[7] The resulting dataset contains 118,754 firm-year observations over 2013–2022 (approximately 12,000 per year), of which 5,779 firms appear in the balanced panel.

Table 1 reports summary statistics for the 2013 cross-section of the analysis sample, broken down by the four 2-digit sectors in Taiwan's Standard Industry Classification (SIC). Following Sutton (2001) and Kugler and Verhoogen (2012), we use R&D intensity (i.e., R&D-to-sales ratio) as our main measure of quality investment, treating R&D as the endogenous sunk investment that builds quality capability.

Two features of these sectors are particularly useful for our research purpose. First, quality investment is ubiquitous: 18% of firms report positive R&D spending, and R&D averages 6% of sales, with both figures highest in the two electronics sectors. Second, the firms rely heavily on the domestic network for their inputs. A majority export directly and exports account for about a quarter of sales, but imports are only about 12% of material costs, and the average firm buys from roughly 200 domestic suppliers. Figure 1 shows that this asymmetry holds throughout 2013–2022 and among the largest firms: the sales-weighted export share exceeds one half, while the sales-weighted import share stays around a quarter, consistent with these

6. In machine tools, for example, more than 98% of Taiwanese producers have fewer than 50 employees (Ministry of Economic Affairs 2005), and more than 70 interviews with makers, suppliers, and industry institutions in the central Taiwan cluster, which hosts over 60% of the industry, find that makers "outsource each step of the production process" to specialized local subcontractors (Chen 2014).

7. Specifically, we exclude firms with annual sales below 1.2 million New Taiwan Dollars (NTD), as these firms are not subject to the standard value-added tax rules. We also exclude firms filing taxes through the "paper review" method; this reflects a long-standing size-dependent tax audit policy in Taiwan, under which firms with annual revenues below 30 million NTD are often exempt from audits by filing taxes through this simplified method. Because filings under this method are rarely audited, the input and output information these firms report in the tax data is less reliable, especially outside of the manufacturing sector. Results in this paper are largely stable even without this selection rule. For reference, 1 US Dollar $\approx$ 30 NTD during the sample period.

industries' large contribution to Taiwan's trade surplus. The vertical chain we study, and the quality investment and rent-sharing along it, are therefore captured largely within the domestic firm-to-firm transaction data.

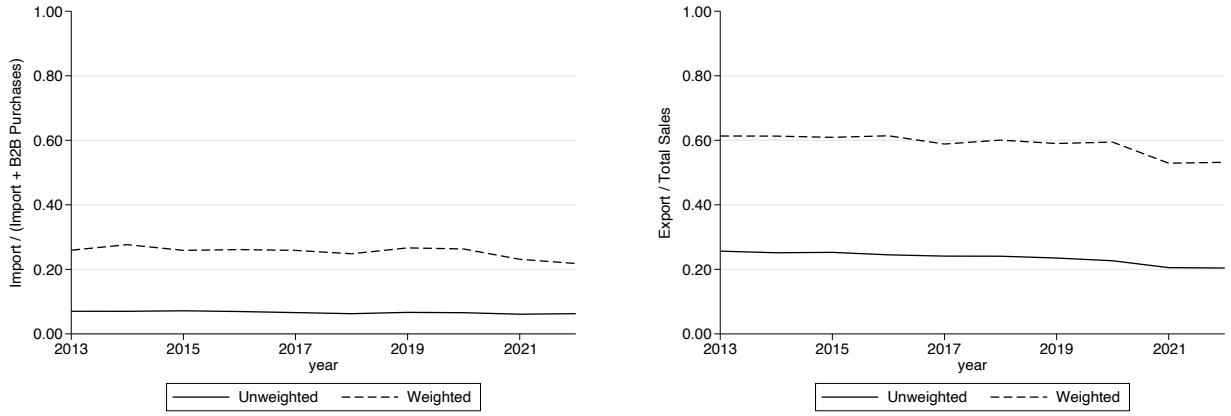
Table 1: Summary Statistics by Sector, 2013

Two-digit SIC code	26 Electronic Parts	27 Computers & Optical	28 Electrical Equipment	29 Machinery & Equipment	All
Total sales (NTD million)	1,210	485	222	134	469
	(12,300)	(2,570)	(1,330)	(635)	(6,310)
Value added (NTD million)	424	134	63	44	156
	(4,560)	(757)	(915)	(328)	(2,350)
Employment	169.7	102.6	45.9	32.0	78.9
	(1,136.5)	(403.9)	(211.1)	(134.4)	(606.2)
1{R&D>0}	0.303	0.332	0.157	0.082	0.185
	(0.460)	(0.471)	(0.364)	(0.275)	(0.389)
1{R&D + advertising>0}	0.716	0.741	0.668	0.668	0.690
	(0.451)	(0.438)	(0.471)	(0.471)	(0.463)
R&D/Sales	0.111	0.118	0.042	0.021	0.061
	(1.307)	(0.740)	(0.525)	(0.725)	(0.887)
(R&D + advertising)/Sales	0.114	0.127	0.046	0.024	0.065
	(1.309)	(0.748)	(0.528)	(0.726)	(0.889)
1{Export>0}	0.643	0.675	0.517	0.564	0.591
	(0.479)	(0.469)	(0.500)	(0.496)	(0.492)
1{Import>0}	0.615	0.666	0.480	0.437	0.521
	(0.487)	(0.472)	(0.500)	(0.496)	(0.500)
Export/Sales	0.277	0.317	0.188	0.240	0.251
	(0.338)	(0.361)	(0.304)	(0.322)	(0.331)
Import/Material Costs	0.166	0.162	0.110	0.071	0.115
	(0.250)	(0.243)	(0.214)	(0.168)	(0.215)
# Suppliers	212.0	218.1	174.2	192.6	197.8
	(283.4)	(257.7)	(188.7)	(184.4)	(225.0)
# Customers	66.2	75.0	93.4	78.6	77.5
	(175.3)	(176.4)	(350.9)	(162.2)	(211.7)
Material/COGS	0.511	0.595	0.671	0.615	0.597
	(0.286)	(0.280)	(0.289)	(0.231)	(0.267)
Observations	2,886	1,578	1,968	4,997	11,429

Note: Means with standard deviations in parentheses, computed on the 2013 cross-section of the analysis sample. # Suppliers and # Customers count domestic B2B partners.

2.2 Measuring Firm-Level Vertical Specialization

To measure firm-level vertical specialization, we first compute upstreamness for each 6-digit SIC industry, following Antràs et al. (2012). After harmonizing industry codes over the sample period, we identify 1,547 6-digit industries in Taiwan, of which 547 belong to manufactur-



(a) Import share of input purchases

(b) Export share of sales

Figure 1: Import and Export Shares of High-Technology Manufacturers, 2013–2022

Note: Panel (a) plots the mean ratio of imports to total input purchases (imports plus domestic B2B purchases); panel (b) plots the mean ratio of exports to total sales. Solid lines are unweighted means across firms; dashed lines are sales-weighted means.

ing. Using the universe of domestic firm-to-firm transactions, before the cleaning and sample restrictions described above, we construct an annual input–output (IO) table at the 6-digit industry level.

We then construct the inter-industry sales-share matrix D, where each entry d_{ij} denotes net sales from industry i to industry j as a share of total sales of industry i. The vector of industry upstreamness is

$$\mathbf{U} = (I - D)^{-1}\,\mathbf{1}$$

where $(I-D)^{-1}$ is the Leontief inverse and $\mathbf{1}$ is a vector of ones. Each entry gives the average number of production stages between an industry and final demand (Antràs et al. 2012). The values are stable throughout our sample period. Energy-related industries are the most upstream, with values around 4, while service industries (including political organizations, religious groups, and betel nut retailers) are the most downstream, with values close to 1, indicating that they sell directly to final demand.

For each firm i, we then compute the transaction-share-weighted upstreamness gap between the firm's suppliers and its own industry:

$$\text{verticalSpan}_{it} = \sum_{j \in \mathcal{J}_{it}^{S}} \frac{x_{ijt}}{\sum_{j' \in \mathcal{J}_{it}^{S}} x_{ij't}} \Big[\text{Upstreamness}(k_{jt}) - \text{Upstreamness}(k_{it})\Big]$$

where $\mathcal{J}_{it}^{S}$ is the set of firm i's suppliers in year t, x_{ijt} the transaction value from supplier j to firm i in the cleaned network, and k_{it} denotes firm i's 6-digit industry in year t. This measure

follows Boehm and Oberfield (2020, 2023) and rests on a similar assumption: controlling for firm size and age, as we will in the later empirical analysis, firms in the same narrowly defined output industry compete in the same product market with the same technology, but can produce with different "recipes", which differ in how processed the inputs they use are. To interpret our later empirical results more intuitively, we define specialization as the negative of vertical span, standardized within each 6-digit-industry-year cell:

$$\text{Specialization}_{it} \equiv -\text{verticalSpan}_{it}$$

A more specialized firm buys inputs closer to its own stage of production and outsources more of the upstream tasks.

Figure 2 illustrates the measure by comparing, within the Metal Machine Tool industry (SIC6 291212, upstreamness = 1.76), firms whose specialization is below the industry's 10th percentile to those above its 90th percentile. We compare their three largest supplier industries. The less specialized firms, shown in the left part of the diagram, source from far-upstream industries: steel castings (241212, upstreamness ≈ 3.0), structural steel (241313, upstreamness ≈ 2.8), and other metal products (259999, upstreamness ≈ 2.0). Presumably, they perform more sequential production steps in-house. By contrast, the more specialized firms, shown on the right, buy over 70% of their domestic inputs from a few supplier industries close to their own: CNC sub-assemblies (291211), precision spindles (291299), and machine tool parts (292999), all with upstreamness close to 1.7.

It is worth mentioning that our specialization measure differs from other measures in the literature based on common ownership (Atalay, Hortaçsu, and Syverson 2014; Atalay et al. 2019). In our data, we can identify purchases from a firm's own branches, as each business entity files VAT returns under its own tax ID and the business registry records each entity's head office. Section A.6 shows that these purchases are negligible: throughout our sample period, about 0.01% of buyer-seller pairs are within-firm, and they account for less than 1% of the sample firms' purchases on average. The own-branch share of purchases is also not significantly related to firm size, age, R&D intensity, or our measure of specialization (Table A8).

In Taiwan, manufacturing firms in these industries rarely own more than one business entity, so few make inputs in separately registered branches. As a result, our measure of specialization captures the make-or-buy choice better than an ownership-based measure would in this context. Such an organizational structure can reflect a scale effect: having a supplier stay independent and serve many buyers can be more cost-efficient than integrating it into one buyer. When input production requires a large fixed investment, an independent supplier spreads that cost over sales to many buyers. Taiwan's machine-tool cluster shows this pattern.

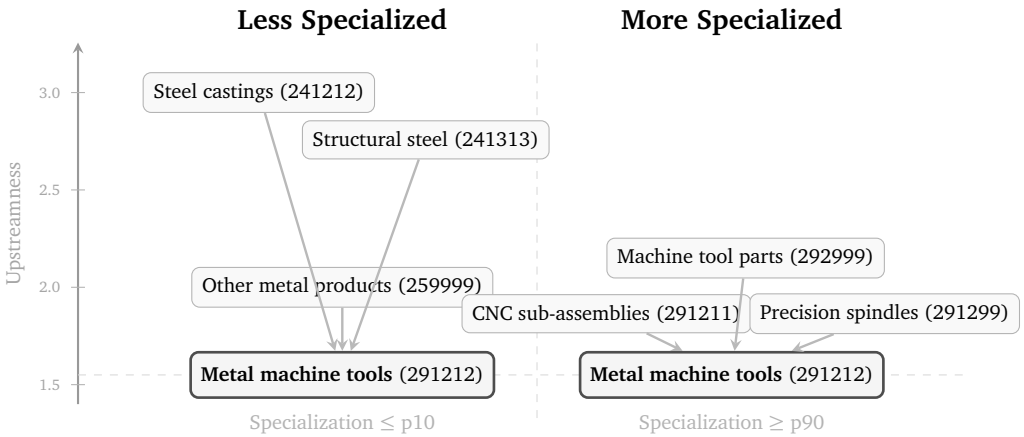


Figure 2: Vertical Specialization in the Metal Machine Tool Industry (SIC6 291212)

Note: Arrows indicate the direction of goods flows from supplier industries to firms in the focal industry. Each panel aggregates the domestic sourcing of a group of firms in the same 6-digit industry (Metal Machine Tools, SIC6 291212, upstreamness = 1.76). The less specialized group (left), with specialization below the 10th percentile of the industry, sources from far-upstream raw material industries; the more specialized group (right), with specialization above the 90th percentile, concentrates sourcing among machine-tool-adjacent suppliers with upstreamness close to that of the focal industry. Numbers in parentheses are 6-digit Standard Industry Classification (SIC6) codes.

One of its main spindle suppliers, for example, began as the spindle division of a machine-tool maker before becoming an independent firm that sells to many makers (Chen 2014). In our model (Section 4), a supplier also spreads the fixed cost of improving input quality over sales to many buyers, while a firm producing in-house bears that cost for its own use alone.

3 Empirical Motivations

We document two facts that link vertical specialization to quality. More specialized firms produce higher-quality output and invest more in quality upgrading (Fact 1). They also buy from higher-quality suppliers, consistent with complementarity between input and output quality (Fact 2). This suggests that a firm's demand for input quality is part of its make-or-buy decision.

Fact 1. Specialized Firms Have Higher Quality

First, using the customs records linked to our data, we measure output quality following Khandelwal (2010). Under constant-elasticity-of-substitution (CES) demand, in which quality enters as a taste shifter, a variety that sells more at a given price is of higher quality. For each

firm i, HS6 product h, destination country c, and year t in the customs data, we estimate the residual of the following demand equation, $\widehat{\xi}_{ihct}$, holding σ fixed, as our measure of log quality:

$$\log x_{ihct} = \delta_{hct} - \sigma \log p_{ihct} + \xi_{ihct},$$

where x_{ihct} is export quantity, p_{ihct} the unit value, δ_{hct} an HS6-product×destination×year fixed effect that absorbs the market's price index and expenditure, and σ the elasticity of substitution, which we set to 5 following Broda and Weinstein (2006).[8] We aggregate the residuals to the firm-year level using export-value weights across HS6 product-destination pairs. Because unit values are observed only in customs records, the quality measure is available only for exporters, who make up more than 60% of the firms in our analysis sample (Table 1).

Figure 3 plots the quality measure against Specialization$_{it}$ among exporters, residualized on industry-by-year fixed effects and firm controls (log sales, log employment, and age). Throughout this section, a positive relationship with Specialization$_{it}$ means that more specialized firms have higher values of the outcome, and all estimates compare firms within the same 6-digit industry and year. Measured output quality rises monotonically with specialization: a one-standard-deviation increase in specialization is associated with a 0.155 log-point increase in output quality (Table 3, column 1). Consistent with this, specialized firms export at higher unit values and pay higher unit values for their imports, within narrowly defined product-country markets (Appendix A).

One potential source of bias in our quality measure is that large exporters may sell at lower unit prices due to quantity discounts. Whether this introduces bias depends on the underlying cause of the discounts. Ganapati and Hottman (2026) document this phenomenon in U.S. import transactions and attribute it to lower per-unit shipping and handling costs. A cost-based discount moves a firm along its demand curve and does not bias the quality measure. Instead, bias would arise if the true demand elasticity differs systematically from the imposed σ across firm sizes. For example, if large Taiwanese exporters dominate their markets and thus face less elastic demand than the imposed σ, our measure would understate their quality, since the imposed elasticity attributes too much of their sales to their low prices and too little to quality. To the extent that vertically specialized firms are larger exporters, our estimated relationship between specialization and quality is therefore conservative. Finally, controlling

8. Under the CES demand assumption, consumers in destination market (h, c, t) allocate expenditure across varieties according to $x_{ihct} = \lambda_{ihct}\, p_{ihct}^{-\sigma}\, P_{hct}^{\sigma-1}\, E_{hct}$, where λ_{ihct} is a demand shifter that we interpret as the quality of firm i's variety, P_{hct} is the quality-adjusted price index, and E_{hct} is market expenditure. Taking logs, the market-level terms $(\sigma-1)\log P_{hct} + \log E_{hct}$ are absorbed by the product-destination-year fixed effect δ_{hct}, so the residual identifies $\log \lambda_{ihct}$: among varieties with the same price in the same market, the one that sells more embodies higher quality. The figures and tables below report the residual divided by $\sigma - 1$, following Khandelwal, Schott, and Wei (2013), who write the shifter as $\lambda^{\sigma-1}$. This constant rescaling leaves the sign and significance of every estimate unchanged.

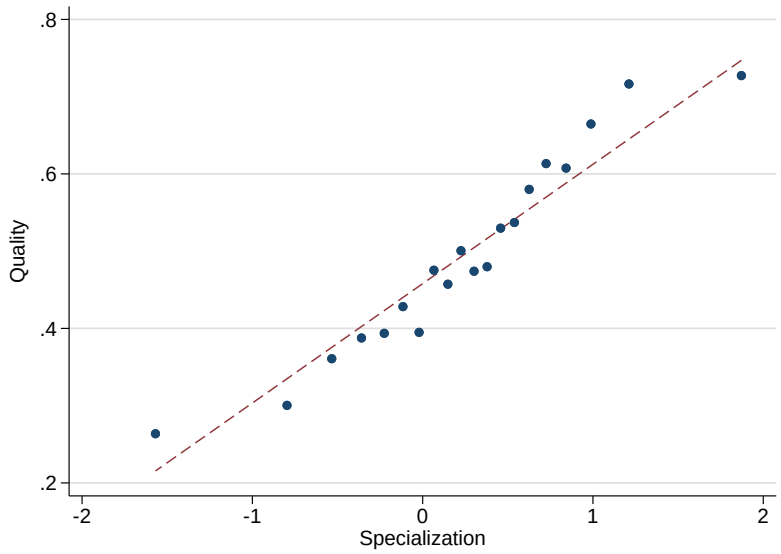


Figure 3: Output Quality and Vertical Specialization

Note: Binscatter of firm-level output quality (export demand residual) on Specialization$_{it}$ among exporters, residualized on SIC6×year fixed effects and firm controls (log sales, log employment, age).

for firm size and age should further mitigate this concern.

Second, corporate income tax records show that specialized firms spend more on quality upgrading. Following Sutton (2001) and Kugler and Verhoogen (2012), we measure quality investment by R&D intensity (RDI, the R&D-to-sales ratio), a broader quality-upgrading intensity that adds advertising (QUI), and intangible-asset intensity (IAI), at both the extensive margin (any positive spending) and the intensive margin (log intensity). We also report the average wage as a proxy for workforce quality. For each outcome y_{it}, we estimate

$$y_{it} = \beta \cdot \text{Specialization}_{it} + \mathbf{X}'_{it}\delta + \mu_{k(i)t} + \varepsilon_{it}$$

where $\mathbf{X}_{it}$ contains log sales, log employment, and firm age, and $\mu_{k(i)t}$ is a SIC6-by-year fixed effect, where $k(i)$ is firm i's 6-digit industry. Table 2 shows that more specialized firms pay higher wages, are more likely to spend on R&D and advertising, to hold intangible assets, and spend/hold more relative to sales on each.

Third, firms' quality-related spending is positively correlated with their measured output quality. Among exporters, we regress the quality measure on each spending measure:

$$\widehat{\text{quality}}_{it} = \beta \cdot y_{it} + \mathbf{X}'_{it}\delta + \mu_{k(i)t} + \varepsilon_{it}$$

Table 2: Vertical Specialization and Quality-Related Expenditure

	Dependent variable y_{it}						
		Any spending, $\mathbf{1}\{\cdot > 0\}$			log spending/sales		
	log wage	R&D	R&D+adv.	Intangibles	R&D	R&D+adv.	Intangibles
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Specialization$_{it}$	0.032	0.034	0.035	0.015	0.371	0.386	0.111
	(0.003)	(0.002)	(0.003)	(0.002)	(0.033)	(0.024)	(0.047)
Controls	log sales, log employment, firm age						
SIC6×year FE	✓	✓	✓	✓	✓	✓	✓
Observations	118,579	118,579	118,579	118,579	19,649	75,879	16,051
R^2	0.294	0.305	0.231	0.207	0.396	0.219	0.277

Note: Each column is a separate regression. Specialization is the negative of vertical span, standardized within industry-year. adv. = advertising; Intangibles = intangible assets. Standard errors in parentheses are clustered at the firm level.

where y_{it} is the spending measure named in each column, and $\mathbf{X}_{it}$ and $\mu_{k(i)t}$ are as above. Column 1 of Table 3 reports the specialization regression behind Figure 3; columns 2–5 show that firms that invest in R&D produce products that command higher demand at a given price.

Taken together, specialized firms produce higher-quality output and spend more on quality-related inputs to production. To our knowledge, this is the first micro-level evidence linking vertical specialization to output quality in advanced manufacturing.

Fact 2. Specialized Firms Sort Toward Higher-Quality Suppliers

Given that vertically specialized firms produce higher-quality goods and outsource a larger share of their upstream tasks, their choice of suppliers is unlikely to be random. Specifically, if these firms purchase inputs mostly from suppliers that also invest heavily in quality, this suggests input-output quality complementarity. This complementarity implies that the marginal return to high-quality inputs is greater for firms producing high-quality outputs, which creates a systematic link between a firm's quality investments, its choice of suppliers, and its degree of vertical specialization.

Figure 4 illustrates the buyer-seller sorting pattern with respect to quality investment in the 2017 cross-section. For each firm, we measure quality-upgrading spending as log (R&D + advertising), residualized on total sales and 6-digit-industry fixed effects, and compute the transaction-share-weighted mean of the same measure across its suppliers. The figure plots this supplier average against the buyer's own spending, separately for the most specialized buyers (specialization at or above the 90th percentile, dashed line) and the least specialized (below the 10th percentile, solid line). Buyers that spend more on quality upgrading buy

Table 3: Quality Spending Maps to Measured Output Quality

	Dependent variable: output quality				
	(1)	(2)	(3)	(4)	(5)
Specialization$_{it}$	0.155				
	(0.012)				
$\mathbf{1}\{\text{R\&D}>0\}_{it}$		0.201			
		(0.025)			
log (R&D/sales)$_{it}$			0.038		
			(0.009)		
$\mathbf{1}\{\text{R\&D+adv.}>0\}_{it}$				0.011	
				(0.020)	
log (R&D+adv.)/sales$_{it}$					0.035
					(0.004)
log Sales$_{it}$	0.094	0.107	0.159	0.103	0.131
	(0.011)	(0.011)	(0.021)	(0.011)	(0.013)
log Emp$_{it}$	0.085	0.042	0.039	0.073	0.046
	(0.013)	(0.014)	(0.028)	(0.014)	(0.016)
Age$_{it}$	−0.009	−0.010	−0.014	−0.011	−0.010
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
SIC6×year FE	✓	✓	✓	✓	✓
Observations	64,484	64,484	17,196	64,484	50,751
R^2	0.173	0.166	0.211	0.162	0.187

Note: Exporters only. Firm-level output quality is the demand residual aggregated by export-value weights. Each column is a separate regression. adv. = advertising. Standard errors in parentheses are clustered at the firm level.

from suppliers that also spend more, and among the most specialized buyers this relationship steepens at high levels of own spending. The rest of this section measures supplier quality by R&D intensity, which we have shown to be strongly and positively correlated with a firm's product quality, and controls for firm size, age, and industry-by-year fixed effects.

To quantify this pattern, we start by calculating the transaction-share-weighted supplier R&D intensity (R&D-to-sales ratio) for each firm:

$$\text{RDI}^s_{it} = \sum_{j \in \mathcal{J}^s_{it}} s^j_{it} \cdot \text{RDI}_{jt}$$

where $\mathcal{J}^s_{it}$ denotes the set of firm i's suppliers in year t and s^j_{it} is firm i's transaction share of supplier j. Furthermore, following Demir et al. (2024), we decompose the weighted supplier RDI into an unweighted mean (Avg) and a covariance (Cov) term:

$$\text{Avg}^s_{it} = \sum_{j \in \mathcal{J}^s_{it}} \frac{1}{|\mathcal{J}^s_{it}|} \text{RDI}_{jt}$$

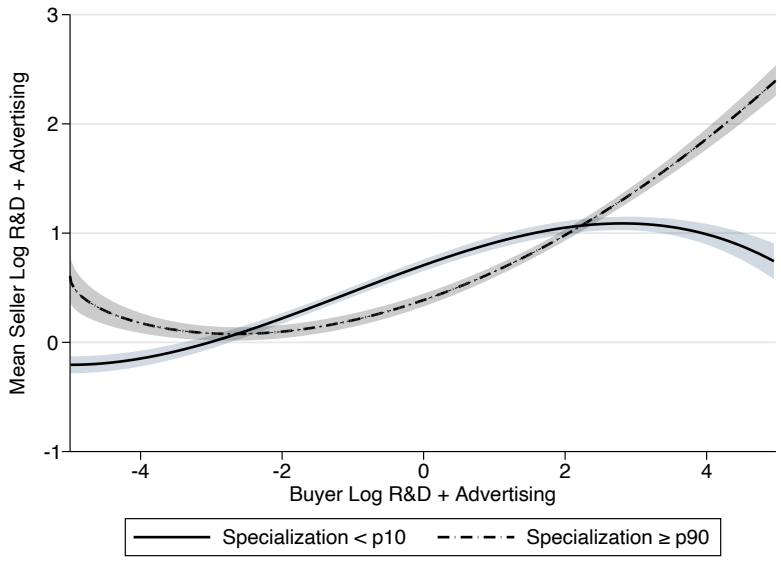


Figure 4: Supplier–Buyer R&D Sorting by Vertical Specialization

Note: The data are from the 2017 cross-section. For both buyers and sellers, quality-upgrading spending is measured as log (R&D + advertising) expenditure, residualized on total sales and 6-digit-industry fixed effects and winsorized at the 1st and 99th percentiles. The vertical axis is, for each buyer, the transaction-share-weighted mean of its suppliers' residualized measure. Curves are fractional-polynomial fits: the dashed line pools the most specialized buyers (specialization at or above the 90th percentile), and the solid line the least specialized (below the 10th percentile). Shaded bands are 95% confidence intervals.

$$\text{Cov}^s_{it} = \text{RDI}^s_{it} - \text{Avg}^s_{it} = \sum_{j \in \mathcal{J}^s_{it}} \left(s^j_{it} - \frac{1}{|\mathcal{J}^s_{it}|} \right) \left(\text{RDI}_{jt} - \text{Avg}^s_{it} \right)$$

Avg captures *whom* the firm buys from, while Cov captures whether *spending tilts toward* higher-quality suppliers, that is, the covariance between transaction shares and supplier R&D intensity. For each supplier-quality measure on the left-hand side, we estimate

$$\begin{aligned} \text{Supplier RDI}^s_{it} \;=\; & \alpha \cdot \text{RDI}_{it} + \beta_1 \cdot \text{Specialization}_{it} \\ & + \beta_2 \cdot \left[\text{RDI}_{it} \times \text{Specialization}_{it} \right] + \mathbf{X}'_{it}\delta + \mu_{k(i)t} + \varepsilon_{it} \end{aligned}$$

where $\mathbf{X}_{it}$ contains log sales, log employment, and firm age, and $\mu_{k(i)t}$ is a SIC6-by-year fixed effect.

Table 4 reports the estimates. We highlight two patterns. First, R&D-intensive firms source from R&D-intensive suppliers, both by choosing suppliers with higher average R&D intensity (Avg) and by concentrating purchases toward the more R&D-intensive ones (Cov). The covariance term accounts for most of this relationship: of the total coefficient on own RDI (0.0046),

the concentration of purchases toward high-RDI suppliers accounts for 0.0040, more than six times the contribution of the average supplier pool (0.0006). Second, holding own R&D intensity fixed, more specialized firms source from more R&D-intensive suppliers. Again, the covariance term accounts for most of the relationship, suggesting that specialized firms allocate a larger share of their purchases to suppliers that invest more in R&D. The positive interaction between own R&D intensity and specialization further suggests stronger sorting on quality investment among specialized firms.[9]

To illustrate the magnitudes, consider a one-standard-deviation increase in specialization, all else equal. Through the level effect, the firm's transaction-share-weighted supplier R&D intensity rises by 0.3 percentage points at any level of its own R&D, which is nearly half the sample mean of 0.7%. Through the interaction, the sorting gradient steepens by roughly 60%, from 0.0046 to 0.0073: for a firm with an R&D-to-sales ratio of 10%, the predicted supplier-quality premium relative to an otherwise identical firm without R&D rises from 0.05 to 0.07 percentage points, i.e., from about 7% to 10% of the sample mean.

The results suggest an economically meaningful role for input-output quality complementarity, as specialized firms match with higher-quality suppliers at every level of own quality. This pattern is consistent with the mechanism we model in Section 4: a firm that specializes in its core stage must secure input quality through the market, so its supplier portfolio carries its quality requirements.

Other Characteristics of Vertically Specialized Firms

We further provide other useful empirical facts in the appendix. We ask which firm characteristics predict specialization when they are considered jointly. Table A3 reverses the regressions of Table 2: it regresses $\text{Specialization}_{it}$ on each quality-related characteristic, controlling for log sales, log employment, firm age, and log wage, with SIC6-by-year fixed effects. Conditional on these characteristics, specialized firms have fewer employees for a given level of sales, consistent with performing fewer production stages in-house. They are also younger and pay higher wages, they are more likely to invest in R&D, advertising, and intangible assets, and they spend more relative to sales on R&D and advertising.

Specialized firms also differ in network-related characteristics. Conditional on size, age, and industry-by-year fixed effects, a one-standard-deviation increase in specialization is associated with about four more suppliers and a 0.03 lower Herfindahl index of purchases, while the customer-side measures, in fact, show no relationship with specialization (Table A4). Specialized firms also keep their supplier links longer, by about 0.13 years on average, and their

9. In Section A.2, we repeat the exercise measuring quality by an indicator for positive R&D and by log RDI (Table A2).

Table 4: Vertical Specialization and R&D Sorting in the Network

	Supplier RDI_{it}		
	Weighted	Avg	Cov
RDI_{it}	0.0046	0.0006	0.0040
	(0.0014)	(0.0003)	(0.0013)
$Specialization_{it}$	0.0033	0.0008	0.0025
	(0.0004)	(0.0003)	(0.0004)
$RDI_{it} \times Specialization_{it}$	0.0027	0.0003	0.0024
	(0.0007)	(0.0002)	(0.0007)
SIC6×year FE	✓	✓	✓
Controls	log sales, log employment, firm age		
Observations	116,488		
R^2	0.042	0.001	0.001

Note: RDI = R&D-to-sales ratio. Weighted = transaction-share-weighted supplier RDI; Avg = unweighted mean supplier RDI; Cov = Weighted − Avg, the covariance between transaction shares and supplier RDI. Controls: log sales, firm age, log employment; SIC6×year fixed effects. Standard errors in parentheses are clustered at the firm level.

supplier links are more likely to survive from one year to the next (Table A5). This stability is consistent with evidence that long-term relationships support quality upgrading (Macchiavello et al. 2026).

Taking stock, the two major facts presented in this section discipline the model that follows. More specialized firms produce higher-quality output and invest more in quality (Fact 1), and they sort toward higher-quality suppliers, so quality complementarities run along the chain (Fact 2). In Appendix A we report further evidence: specialized firms export at higher unit values and pay higher unit values for their imports, and, following a shift-share export demand shock, they tilt their sourcing toward more R&D-intensive suppliers, though we treat this last result as suggestive.

These facts also point to a cost of using the market that falls disproportionately on specialized firms. Because these firms buy more of their inputs from other firms, they pay costs that in-house production avoids, most directly the markups their suppliers charge. The model in the next section makes these markups explicit and asks how they shape firms' outsourcing and quality investment.

4 The Network Model

The model features sequential production as in Yi (2003) and Johnson and Moxnes (2023), with input-output quality complementarities as in Fieler, Eslava, and Xu (2018) and Demir et al. (2024). Consider a bipartite network economy with two manufacturing sectors, one upstream and one downstream. A representative consumer combines downstream varieties and supplies L units of labor inelastically at the wage w.

Each of the two manufacturing sectors has a unit measure of potential firms producing a continuum of varieties on the unit interval. We denote the sets of upstream and downstream varieties by $\mathscr{U}$ and $\mathscr{D}$, and use $u \in \mathscr{U}$ to index upstream firms and $d \in \mathscr{D}$ to index downstream firms. We use h to index a downstream firm's in-house department that provides inputs within the firm's boundary. Entry works the same way in both sectors. The measure of potential entrants is exogenous, each potential entrant observes its own draws of quality-upgrading ability, and it becomes active only if operating at those draws yields positive profit. The measures of active upstream and downstream varieties are therefore both equilibrium outcomes.

Figure 5 summarizes the environment. A downstream firm, given its abilities to upgrade the quality of its output (ω_1^d) and of its in-house inputs (ω_1^h), assembles its material input from two sources: it buys from external upstream suppliers of heterogeneous quality-upgrading ability ω_1^u, paying their markup but gaining access to their quality, and it produces the input in-house at marginal cost, avoiding the markup. The downstream firm sells its output to domestic consumers and, once we open the economy in Section 4.2, to the export market. The make-or-buy decision, the quality investments of upstream and downstream firms, and the complementarity between input and output quality are the objects of the analysis.

In our model, every active supplier sells to every downstream firm, and every downstream firm buys from every active supplier. What varies is how much each buyer spends on each supplier, that is, the intensive margin of the network. Motivated by Fact 2, in which the Cov term in Table 4 accounts for most of the buyer-seller sorting, we let this expenditure share depend on the supplier's quality and on the buyer's output quality. On top of comparative advantage, the strength of input-output quality complementarity governs how strongly the share responds to these qualities.

The rest of this section proceeds as follows. Section 4.1 sets out the closed economy, defines its equilibrium, and derives the link between quality investment and vertical specialization. Section 4.2 opens the economy to exports. Section 4.3 compares the decentralized allocation with the planner's to show how suppliers' markups distort outsourcing and quality choices. Section 4.4 shows how a firm's exposure to suppliers' markups varies with its output quality through how much it outsources.

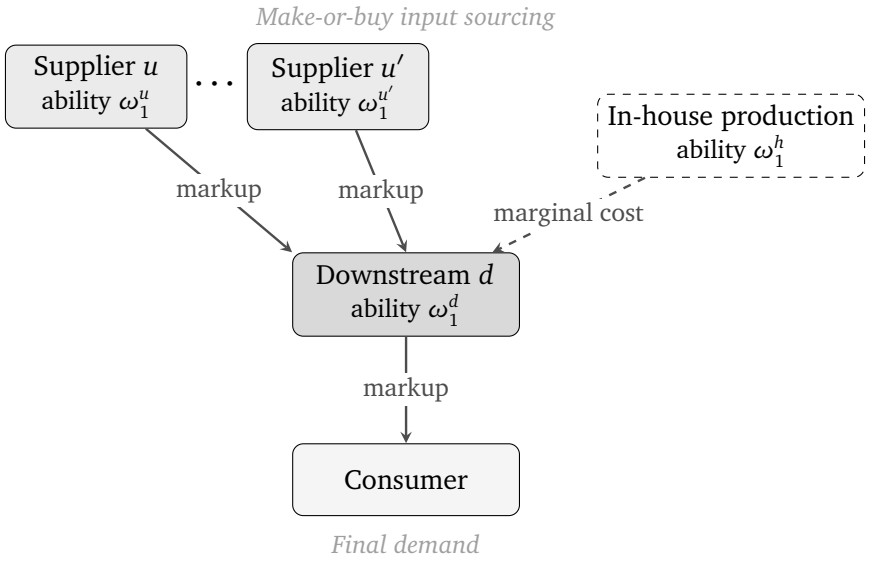


Figure 5: Structure of the Network Model

Note: Arrows show the flow of inputs. A downstream firm buys from a continuum of external suppliers (the ellipsis) at prices that include a markup, or makes the input in-house at marginal cost (dashed box). In-house quality is limited by the firm's own upstream ability ω_1^h; each ω_1 is the quality-upgrading ability. The figure shows the closed economy. Section 4.2 adds an export market for downstream output.

4.1 The Closed Economy

The Upstream Firm's Problem. In the upstream sector, a firm $u \in \mathscr{U}$ draws $\omega^u \equiv (\omega_0^u, \omega_1^u)$, a Hicks-neutral productivity shifter and a quality-upgrading ability. The ability ω_1^u governs the return to quality investment, i.e., how much the firm's productivity rises with the quality it chooses. It produces a unique input variety of quality q^u with a labor-only technology:[10]

$$y^{\mathscr{U}}(u, q^u) = z(\omega^u, q^u)\,\ell \tag{1}$$

where $z(\omega^u, q^u)$ is firm u's productivity.

As a supplier, the firm charges a constant markup over its marginal cost:

$$p^{\mathscr{U}}(u) = \mu \frac{w}{z(\omega^u,\, q^u)} \tag{2}$$

where $\mu \equiv \sigma/(\sigma - 1)$. It sells its variety to downstream firms, which combine it with other inputs in their material bundles as described below. Because it sells to every downstream firm,

10. Upstream production, whether by a supplier or in-house, uses labor only with constant returns to scale, so its marginal cost is w/z and all of its variable cost is labor. This keeps the upstream sector intentionally simple: upstream firms represent local, labor-intensive production.

the supplier spreads the fixed cost of quality upgrading over all of its buyers. This is the scale advantage of independent suppliers discussed in Section 2. Its operating revenue is

$$r^{\mathscr{U}}(u) \equiv r^{\mathscr{U}}(q^u, p^{\mathscr{U}}) = \left(p^{\mathscr{U}}\right)^{1-\sigma} D^{\mathscr{U}}(q^u)$$

where $D^{\mathscr{U}}(q^u)$ is the demand shifter the supplier faces from downstream buyers. Active upstream firms maximize profits by choosing quality:

$$\pi(u) = \max_{q^u} \sigma^{-1} r^{\mathscr{U}}(q^u, p^{\mathscr{U}}) - w f(q^u) \tag{3}$$

where the first term is the operating profit and the second is the fixed cost of quality upgrading $f(q^u)$, denominated in units of labor and specified below.

A potential upstream firm observes its draws and becomes active only if the maximized profit at those draws is non-negative, and otherwise stays out of the market. The measure of active upstream firms, and hence the set of external input varieties available to downstream firms, is determined in equilibrium. For an active firm choosing an interior quality, $q^u > 0$, the first-order condition is:

$$\sigma^{-1} \frac{\partial r^{\mathscr{U}}(q^u, p^{\mathscr{U}})}{\partial q^u} = w f'(q^u) \tag{4}$$

The condition equates the marginal operating profit from quality with its marginal upgrading cost. The left-hand side also shows the role of the markup. The supplier keeps a share $\sigma^{-1} = 1 - 1/\mu$ of its revenue as operating profit, and this profit is what pays for quality. If the supplier had to sell at marginal cost, this profit would be zero at every quality, so it would choose $q^u = 0$ and leave the market. Downstream firms would then make more of their inputs in-house, with consequences for output quality and real consumption. The markup thus raises input prices but also finances upstream quality upgrading. Nothing here hinges on the markup being a constant proportion of cost. Any arrangement under which suppliers keep part of the surplus from firm-to-firm trade can pay for their quality investment, provided that both upstream and downstream quality matter for final output. Section 6 quantifies both sides of this trade-off and discusses its implications for industrial policy.

The Downstream Firm's Production. In the downstream sector, a firm $d \in \mathscr{D}$ draws $\omega^d \equiv (\omega_0^d, \omega_1^d)$ for its core production and $\omega^h \equiv (\omega_0^h, \omega_1^h)$ for its in-house production. The ability ω_1^d governs the return to quality upgrading in core production, and ω_1^h governs it in in-house production. Their correlation, $\varrho_\omega \equiv \text{corr}(\omega_1^d, \omega_1^h)$, governs the firm's comparative advantage and its gains from specialization. The Hicks-neutral shocks, ω_0^d and ω_0^h, are drawn independently of each other and of the abilities. The firm can perform the upstream task in-house with the

same constant-returns, labor-only technology as an upstream firm, producing $y^h(d)$ at chosen quality q^h:

$$y^h(d, q^h) = z(\omega^h, q^h)\,\ell \tag{5}$$

In its core production, given its chosen output quality q^d, the downstream firm combines labor and a material bundle m that draws on two sourcing channels, external upstream varieties and its own in-house variety:

$$\begin{aligned} y^{\mathscr{D}}(d, q^d) &= z(\omega^d, q^d)\,\ell^{\alpha_\ell}\, m^{\alpha_m} \\ m &= \left[\int_{u\in\mathscr{U}} \phi(q^d, q^u)^{1/\sigma}\, x^u(d)^{\frac{\sigma-1}{\sigma}}\, du \;+\; [\eta\phi(q^d, q^h)]^{1/\sigma}\, x^h(d)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \end{aligned} \tag{6}$$

where α_ℓ and α_m are the labor and material cost shares, with $\alpha_\ell + \alpha_m = 1$. The material aggregate m is a CES aggregate of the firm's purchases from upstream firm u, denoted $x^u(d)$, and its use of its own in-house variety, denoted $x^h(d)$. The function ϕ captures the complementarity between the quality of input varieties (q^u for external supplier u, and q^h for the in-house input) and the firm's output quality q^d. The parameter $\eta > 0$ is the CES weight on the in-house variety, which governs the *effective measure* of the in-house material: the aggregator treats the firm as if it produced a mass η of in-house varieties, all of the same chosen quality q^h and each imperfectly substitutable with the external varieties. The elasticity of substitution σ, which also determines the markup $\mu = \sigma/(\sigma-1)$ charged by suppliers, governs how sourcing shares across the two channels respond to relative prices, adjusted by ϕ and η.[11]

Because materials are a CES aggregate, no firm is fully specialized. Throughout, we write q^u, q^d, and q^h for the firm's chosen quality levels at the upstream core, downstream core, and downstream in-house stages. These are equilibrium objects that depend on the firm's ability draws and other firms' choices.

Quality and Productivity. Each production node introduced above (an upstream firm, a downstream firm's core production, and its in-house production) operates at a chosen quality q with its draws $\omega = (\omega_0, \omega_1)$. Following Demir et al. (2024), productivity takes the form

$$z(\omega, q) = \exp\!\big(\omega_0 + \omega_1 \log q + \bar{\omega}(\log q)^2\big)$$

The shifter ω_0 captures a firm's absolute advantage in scale, as it increases productivity at every quality level. It does not change the elasticity of productivity with respect to quality

11. The aggregate m simplifies a multi-industry sourcing technology in which a firm, within each upstream industry, chooses between a continuum of external varieties and an in-house substitute.

and is drawn independently of ω_1, with a separate draw at each production node: ω_0^u for an upstream firm and (ω_0^d, ω_0^h) for a downstream firm's core and in-house production. The independence from ω_1 separates firms that are large because they are efficient from firms that are large because they invest in quality, and allows the model to generate dispersion in sales that is not entirely tied to the quality distribution.[12] Firms differ only through these draws, so every firm-level equilibrium object below is a function of ω^u for an upstream firm and of (ω^d, ω^h) for a downstream firm.

The ability ω_1 captures a firm's comparative advantage in quality, governing the return to upgrading. Quality upgrading is costly: the fixed overhead cost $f(q)$ in the firms' problems above is increasing in q and denominated in units of labor, with $f(0) = 0$. The same cost function applies to every quality choice in the model. Quality also translates into productivity nonlinearly: because the exponent is quadratic in $\log q$, the elasticity of productivity with respect to quality varies with the quality level, and $\bar{\omega}$ governs this nonlinearity. When $\bar{\omega} < 0$, log productivity is concave in log quality: at high quality, raising it further lowers productive efficiency and thus raises marginal cost. In addition to providing the curvature for solving the model, this parameterization choice permits a richer mapping between quality and unit price.

Input-output Quality Complementarity. Following Demir et al. (2024), we parameterize the complementarity function as

$$\phi(q^d, q^u) = \frac{\exp(q^u - \nu q^d)}{1 + \exp(q^u - \nu q^d)}, \quad \nu \geq 0$$

As in Fieler, Eslava, and Xu (2018) and Demir et al. (2024), the parameter ν captures the degree of quality complementarity and is an object of primary interest in the quantitative exercise. The function exhibits the following properties: (1) $\partial\phi/\partial q^u > 0$; (2) $\partial\phi/\partial q^d < 0$; (3) log-supermodularity, $\partial^2 \log\phi/\partial q^u \partial q^d \geq 0$, meaning that the return to raising input quality is relatively larger when output quality is higher.

To see how the strength of quality complementarity affects the make-or-buy margin, holding everything else fixed, note that under the CES material aggregate the sourcing ratio between the external and in-house varieties is

$$\frac{x^u(d)}{x^h(d)} = \left(\frac{p^{\mathscr{U}}(q^u)}{p^h(q^h)}\right)^{-\sigma} \frac{\phi(q^d, q^u)}{\eta\phi(q^d, q^h)} \propto \mu^{-\sigma} \frac{\phi(q^d, q^u)}{\eta\phi(q^d, q^h)}$$

12. A feature of the model that departs from Demir et al. (2024) is that more efficient firms actually tend to choose higher quality, even though ω_0 is drawn independently of ω_1. This is because a firm with a higher ω_0 sells more at any quality, and the same quality upgrade pays off over a larger base of sales, making the fixed cost of quality easier to cover. This scale effect on quality choice corresponds to the correlation between sales and R&D spending in the data, although it is outside the set of targeted moments.

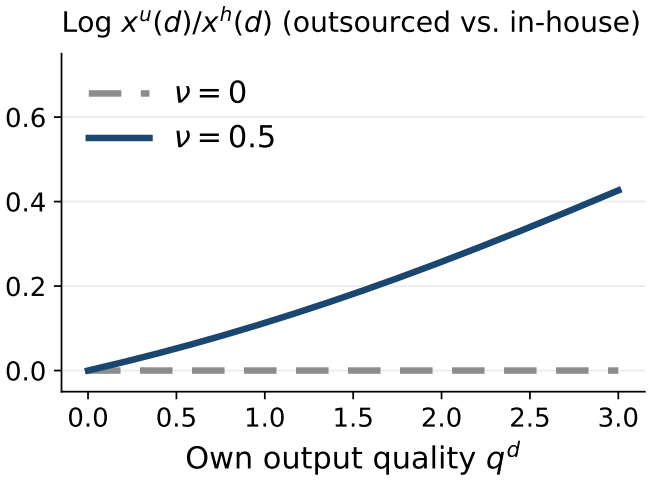


Figure 6: Complementarity Tilts Sourcing Toward Higher-Quality Suppliers

Note: Numerical illustration of Equation 7: the log ratio of purchases from one external supplier of quality $q^u = 2.5$ to use of an in-house variety of quality $q^h = 1$, against the firm's output quality q^d, with prices and input qualities held fixed. The level is normalized to zero at $q^d = 0$. With $\nu = 0$, the ratio does not depend on q^d.

The ratio is therefore increasing in q^d for any $q^u > q^h$: since $\partial \log \phi(q^d, q)/\partial q^d = -\nu[1 - \phi(q^d, q)]$,

$$\frac{\partial}{\partial q^d}\left(\log \frac{x^u(d)}{x^h(d)}\right) = \nu\left[\phi(q^d, q^u) - \phi(q^d, q^h)\right] \geq 0 \tag{7}$$

Figure 6 illustrates this comparative static numerically: it shows how the ratio of outsourced to in-house inputs changes with a downstream firm's output quality, all else equal. With complementarity ($\nu > 0$), a firm that raises its output quality shifts its purchases toward suppliers whose quality exceeds that of its in-house variety.

The parameter ν has a direct economic meaning. The value of $\phi(q^d, q^u)$ depends on the two qualities only through $q^u - \nu q^d$, so ν is the increase in input quality needed to keep an input's weight in the bundle unchanged when the buyer's output quality rises by one unit. When $\nu = 0$, an input's weight does not depend on who buys it, and the sourcing ratio in Figure 6 is flat. When $\nu > 0$, a firm that upgrades its output must also upgrade its inputs to keep them equally useful, and it can do so either by raising q^h in-house or by buying from higher-quality suppliers.

The Downstream Firm's Problem: Pricing and Profit. A downstream firm $d \in \mathscr{D}$ with output quality q^d and in-house upstream quality q^h combines the two sourcing channels into a

material price index:

$$P_m(q^d, q^h) = \left[\int_{u \in \mathscr{U}} \phi(q^d, q^u)\, p^{\mathscr{U}}(u)^{1-\sigma}\, du \;+\; \eta \phi(q^d, q^h)(p^h)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{8}$$

where the in-house variety is priced at marginal cost, $p^h = w/z(\omega^h, q^h)$, since the firm supplies it to itself. External inputs, on the other hand, carry the upstream markup μ. The downstream firm then charges the same constant markup over its own marginal cost when selling to the consumer:

$$p^{\mathscr{D}}(d) = \mu \frac{w^{\alpha_\ell}\, P_m(q^d, q^h)^{\alpha_m}}{z(\omega^d, q^d)} \tag{9}$$

The representative consumer's consumption bundle, the sectoral good y_s, combines all downstream varieties, weighted by quality through $\Phi(q^d) \equiv \phi(0, q^d)$:

$$y_s = \left[\int_{d \in \mathscr{D}} \Phi(q^d(d))^{\frac{1}{\sigma}}\, y^{\mathscr{D}}(d)^{\frac{\sigma-1}{\sigma}}\, dd \right]^{\frac{\sigma}{\sigma-1}} \tag{10}$$

The consumer spends total income Y, consisting of wage income and manufacturing firm profits, on this sectoral final output; its price index is

$$p_s = \left[\int_{d \in \mathscr{D}} \Phi(q^d(d))\, p^{\mathscr{D}}(d)^{1-\sigma}\, dd \right]^{\frac{1}{1-\sigma}} \tag{11}$$

Taking the price index of the sectoral good p_s and aggregate income Y as given, we can express the downstream firm's operating revenue as

$$\begin{aligned} r^{\mathscr{D}}(d) &= \left(p^{\mathscr{D}}\right)^{1-\sigma} D^{\mathscr{D}}(q^d), \\ D^{\mathscr{D}}(q^d) &= \Phi(q^d)\, p_s^{\sigma-1}\, Y \end{aligned} \tag{12}$$

where $D^{\mathscr{D}}(q^d)$ is domestic demand from the consumer. The firm chooses its two qualities jointly:

$$\pi(d) = \max_{q^d, q^h} \; \sigma^{-1}\left(p^{\mathscr{D}}\right)^{1-\sigma} D^{\mathscr{D}}(q^d) \;-\; wf(q^d) - wf(q^h) \tag{13}$$

The two choices interact through the material price index: raising q^d shifts the complementarity weights $\phi(q^d, \cdot)$ across all input varieties, including the in-house one, so the return to in-house quality depends on the output quality the firm has chosen, and vice versa. As in the upstream sector, entry depends on the maximized profit at the firm's own draws: a potential downstream firm becomes active only if $\pi(d)$ is non-negative, so the measure of active downstream varieties,

and hence the set of varieties in the consumer's bundle, is an equilibrium outcome.[13]

The Make-or-Buy Trade-off. For an interior output quality choice, the first-order condition can be written as

$$\underbrace{(\sigma-1)\frac{\partial \ln z(\omega^d, q^d)}{\partial q^d}}_{\text{core efficiency}} + \underbrace{\frac{\partial \ln D^{\mathscr{D}}(q^d)}{\partial q^d}}_{\text{demand}} = \underbrace{\alpha_m(\sigma-1)\frac{\partial \ln P_m}{\partial q^d}}_{\text{input costs}} + \underbrace{\frac{wf'(q^d)}{\sigma^{-1}(p^{\mathscr{D}})^{1-\sigma}D^{\mathscr{D}}(q^d)}}_{\text{cost of the upgrade}} \tag{14}$$

The left-hand side displays how upgrading q^d changes core productivity and demand; the right-hand side summarizes the cost, through a higher material price index and the investment itself, relative to operating profit. Complementarity enters through the input-cost term.

To see how input costs change when output quality increases, we derive the partial derivative using Equation 8:

$$\alpha_m(\sigma-1)\frac{\partial \ln P_m}{\partial q^d} = \alpha_m\, \nu(1-\bar{\phi}), \qquad \bar{\phi} \equiv \int_{\mathscr{U}} s^u(d)\,\phi(q^d, q^u)\,du + s^h(d)\,\phi(q^d, q^h)$$

where $s^u(d) = \phi(q^d, q^u)(p^{\mathscr{U}})^{1-\sigma}/P_m^{1-\sigma}$ and $s^h(d) = \eta\phi(q^d, q^h)(p^h)^{1-\sigma}/P_m^{1-\sigma}$ are the shares of supplier u and of the in-house variety in the firm's material spending. Raising q^d makes the firm's current input bundle less complementary to its output, and more so when ν is larger.

The firm can respond in two ways as it chooses a higher output quality. It can simultaneously invest more in in-house quality q^h, at cost $f(q^h)$, whose return is bounded by its in-house ability ω_1^h. The other choice is to shift its sourcing toward higher-quality suppliers, paying the markup μ on every unit. Which response a firm takes depends mostly on two mechanisms. The first is comparative advantage, governed by the correlation ϱ_ω between the firm's abilities (ω_1^d, ω_1^h): when this correlation is weakly positive or even negative, a firm that is good at upgrading its output is often not good at upgrading its in-house input, so firms that find it worthwhile to raise q^d need not find it worthwhile to raise q^h.[14] The second is complementarity: ν amplifies this gain from outsourcing and shapes which suppliers each firm buys from, as we see in the sorting pattern in Taiwan's network data (Section 3).

13. At the estimated parameters, nearly all potential downstream entrants find it profitable to enter. The downstream entry condition is therefore effectively nonbinding, and the measure of active firms is approximately equal to the pool of potential entrants.

14. An alternative way to model this trade-off is a span-of-control model, in which spreading a firm's fixed ability across more quality-upgrading tasks lowers the return to each.

Upstream Demand. Aggregating the material spending of downstream firms gives the demand shifter that each supplier faces:

$$D^{\mathscr{U}}(q^u) = \frac{\alpha_m}{\mu} \int_{d \in \mathscr{D}} \phi(q^d(d), q^u) \left[P_m(d)\right]^{\sigma-1} r^{\mathscr{D}}(d)\, dd \tag{15}$$

where $r^{\mathscr{D}}(d)$ is the operating revenue of downstream firm d. The factor α_m/μ is material spending as a share of the downstream firm's revenue: the firm's costs are $r^{\mathscr{D}}/\mu$, and it spends a share α_m of them on materials. The demand shifter $D^{\mathscr{U}}(q^u)$ aggregates spending from all downstream manufacturers that source the upstream variety u, and it captures the supplier's return to quality upgrading: raising q^u raises the weight $\phi(q^d, q^u)$ that downstream buyers place on its variety. With complementarity, this gain is larger among high-quality buyers, so a supplier's return to upgrading rises with the quality of the buyers it serves. Higher ability ω_1^u also lowers a supplier's cost through productivity in most cases. Together, the two effects are the supply side of the quality sorting in Section 3: suppliers with better buyers and higher ability choose higher quality. The shifter is an equilibrium object. What other suppliers choose affects downstream firms' sourcing and output quality, and those choices in turn set each supplier's return to quality investment. We solve for the shifter as a fixed point (Appendix B).

Equilibrium. To focus on the product market and firms' decisions about quality spending and specialization, we take the wage as the numeraire and fix it at $w = 1$; the consumer supplies L units of labor inelastically. The labor market then clears not through wage adjustment but through the equilibrium response of the sectoral price index p_s and aggregate income Y.

An equilibrium consists of quality policies $q^u(\omega^u)$, $q^d(\omega^d, \omega^h)$, and $q^h(\omega^d, \omega^h)$, together with the upstream demand shifter $D^{\mathscr{U}}(\cdot)$, the sectoral price index p_s, and aggregate income Y, such that: (1) taking $D^{\mathscr{U}}$, p_s, and Y as given, $q^u(\omega^u)$ maximizes $\pi(u)$ for every ω^u, with the price $p^{\mathscr{U}}$ then set by Equation 2; (2) taking the upstream price schedule and the aggregates as given, $\left(q^d(\omega^d, \omega^h), q^h(\omega^d, \omega^h)\right)$ maximizes $\pi(d)$ for every (ω^d, ω^h); (3) the aggregates are generated by the firms' own choices: integrating the type-contingent policies over the type distributions reproduces the upstream demand shifter $D^{\mathscr{U}}$ in Equation 15 and the sectoral price index p_s in Equation 11; and (4) goods and labor markets clear. Each supplier's output equals the purchases routed to it, $y^{\mathscr{U}}(u) = \int_{\mathscr{D}} x^u(d)\, dd$, and the in-house variety serves only its own firm, $y^h(d) = x^h(d)$. Goods-market clearing equates nominal spending on the sectoral good to total income,

$$Y = wL + \Pi^u + \Pi^d \tag{16}$$

where $\Pi^u \equiv \int_{\mathscr{U}} \pi(u)\, du$ and $\Pi^d \equiv \int_{\mathscr{D}} \pi(d)\, dd$ are aggregate upstream and downstream profits

that are rebated to the consumer as the owner of firms in the economy. Labor-market clearing, as shown in Equation 17 below, allocates the endowment L to production labor at the three tasks (upstream, downstream core, and downstream in-house; first line) and to quality-upgrading overheads at each stage (second line):

$$\begin{aligned} L = &\int_{\mathscr{U}} \ell^u(u)\,du + \int_{\mathscr{D}} \ell^d(d)\,dd + \int_{\mathscr{D}} \ell^h(d)\,dd \\ &+ \int_{\mathscr{U}} f(q^u(u))\,du + \int_{\mathscr{D}} \left[f(q^d(d)) + f(q^h(d))\right] dd \end{aligned} \tag{17}$$

Mechanism: Quality Choice and Vertical Specialization. The model links a downstream firm's output quality, the quality of the suppliers it buys from, and its degree of vertical specialization, consistent with the empirical patterns documented in Section 3. Among firms with similar in-house ability ω_1^h, those with higher core ability ω_1^d invest in higher output quality q^d. As q^d rises, the log-supermodularity of ϕ tilts sourcing toward high-quality external suppliers and away from the in-house variety. Although the in-house option retains a cost advantage (no markup), the fixed cost $f(q^h)$ and the in-house ability ω_1^h limit how far q^h can rise to keep pace with q^d. As a result, the in-house share of material spending is decreasing in output quality q^d among firms with similar ω_1^h: higher-quality downstream firms are more vertically specialized and source inputs disproportionately from higher-quality suppliers.

4.2 The Open Economy

Before taking the model to the data, we add an export decision to the downstream firm's problem. We abstract from imports in the baseline model because the make-or-buy margin plays out largely within the domestic network (Figure 1 and Table 1).[15] By contrast, exports account for about a quarter of sales, or about 60% when weighted by firm sales (Figure 1), as these sectors sit upstream in international supply chains for quality-critical manufacturing goods. Connections with high-performing international buyers encourage domestic downstream firms to invest in quality, which in turn encourages upstream development at home.

In this exercise, downstream firms take world demand as given. Their export revenue depends on the real exchange rate e, the price of foreign goods relative to home goods; e is normalized to one in the estimated equilibrium and adjusts in the counterfactuals, as described in the equilibrium definition below.[16] The core mechanism of the closed economy is unchanged.

15. In the data, the import share of input purchases is highest among the most integrated firms, even though specialized firms are more likely to import (Table D2). Appendix D extends the model with an import margin that speaks to this pattern.

16. Upstream firms, which use a labor-only technology, are local and do not trade internationally; only down-

Foreign Demand. On the demand side, the export revenue of downstream variety d with quality q^d and price $p^{\mathscr{D}}$ is

$$r_F^{\mathscr{D}}(d) = \left(p^{\mathscr{D}}\right)^{1-\sigma} e^{\sigma} D_F(q^d)$$

where $D_F(q^d)$ is an exogenous foreign demand shifter, parameterized as

$$D_F(q^d) = b_1 (q^d)^{b_2}$$

with b_1 governing the level of foreign demand and b_2 how steeply it rises with quality.

After quality choices, firms independently draw the sunk cost for exporting, f_E, denominated in units of labor. This allows firms with similar quality and productivity to differ in export status.

The Downstream Firm's Problem with Exports. The extensive margin of the export decision appears in the downstream firm's problem through a discrete choice, the export status $E \in \{0, 1\}$. Conditional on export status, the export revenue shows up as an additional part of the firm's operating revenue:

$$\begin{aligned} r^{\mathscr{D}}(d, E) &= \left(p^{\mathscr{D}}\right)^{1-\sigma} D^{\mathscr{D}}(q^d, E), \\ D^{\mathscr{D}}(q^d, E) &= \Phi(q^d)\, p_s^{\sigma-1}\, Y + E \cdot e^{\sigma} D_F(q^d) \end{aligned} \tag{18}$$

where $p^{\mathscr{D}}$ is the price in Equation 9.

The firm chooses its qualities before observing the fixed-cost draw. Based on this timing assumption, the updated downstream firm's problem is

$$\pi(d) = \max_{q^d, q^h} \mathbb{E}_{f_E}\left[\max_{E \in \{0,1\}} \sigma^{-1}\left(p^{\mathscr{D}}\right)^{1-\sigma} D^{\mathscr{D}}(q^d, E) \; - \; w E f_E \right] - wf(q^d) - wf(q^h) \tag{19}$$

The inner maximization selects the export status given the realized draw, where the firm exports if and only if the export profit, $\sigma^{-1}(p^{\mathscr{D}})^{1-\sigma} e^{\sigma} D_F(q^d)$, is larger than the fixed cost of exporting, wf_E. The outer expectation is over the distribution of the fixed cost. Higher-quality firms are more likely to export: foreign demand rises with quality when $b_2 > 0$, so the export margin reinforces the incentive to upgrade quality beyond domestic demand.

Equilibrium with Exports. The equilibrium definition of the closed economy extends in three ways. First, the real exchange rate e joins the list of aggregates that firms take as given. Second, the export statuses and sales will affect upstream firms through changes in the upstream

stream firms export.

demand shifter, where the operating revenue $r^{\mathscr{D}}(d)$ in Equation 15 is replaced by $r^{\mathscr{D}}(d, E(d))$. Downstream firms producing at a higher q^d are more likely to export, which in turn generates more upstream demand and incentivizes capable upstream firms to upgrade quality. Third, the export flow enters the market clearing conditions for the sectoral good market and the labor market. To see that, note that aggregate export revenue is

$$X = \int_{\mathscr{D}} E(d)\, r_F^{\mathscr{D}}(d)\, dd \tag{20}$$

As we abstract from importing in our baseline model, export revenue is also the economy's current-account surplus, which is the income earned abroad and not spent at home.[17] Goods-market clearing therefore equates domestic spending on the final good to total income net of exports,

$$Y = wL + \Pi^u + \Pi^d - X \tag{21}$$

and real domestic absorption is $C = Y/p_s$. The labor-market clearing condition shown in Equation 17 now has to consider the export fixed costs:

$$\begin{aligned} L = &\int_{\mathscr{U}} \ell^u(u)\, du + \int_{\mathscr{D}} \ell^d(d)\, dd + \int_{\mathscr{D}} \ell^h(d)\, dd \\ &+ \int_{\mathscr{U}} f(q^u(u))\, du + \int_{\mathscr{D}} \left[f(q^d(d)) + f(q^h(d))\right] dd + \int_{\mathscr{D}} E(d)\, f_E(d)\, dd \end{aligned} \tag{22}$$

With exports entering the model, it is crucial to give proper treatment to the current account when discussing the real domestic absorption, which is the key objective for measuring welfare, in counterfactual analyses. Given the wage rate w and real exchange rate e, the market-clearing conditions determine domestic spending Y. In the estimated equilibrium, $e = 1$. In the policy counterfactuals of Section 6.3, we let e adjust so that real exports, X/p_s, stay at their baseline level: the economy ships abroad the same quantity of its output. This allows us to understand the consequence of a changing quality distribution in domestic consumption rather than in a larger trade surplus, which would divert labor from producing for domestic consumption.[18]

17. We can rationalize this surplus by assuming that all export revenue is held as international assets outside the model. This simplifies the problem and allows us to focus on the quality choices in the estimated equilibrium.

18. In a counterfactual, higher exports are income saved abroad rather than consumption at home, so we hold real exports fixed and measure only the gain at home. When no exchange rate can keep real exports at their baseline level, as when suppliers' markups are removed in Section 6.3, we keep e at one and report the fall in exports.

4.3 The Planner's Problem

We now turn to the planner's problem to characterize how the decentralized equilibrium deviates from the social optimum, shedding light on the quality-markup tension. The planner maximizes the consumer's bundle subject to the same labor resource constraint that clears the decentralized economy, Equation 17:

$$\max_{\{q^d(d),\, q^u(u),\, q^h(d)\}} Y_s = \left[\int_{d\in\mathscr{D}} \Phi(q^d(d))^{\frac{1}{\sigma}}\, y^{\mathscr{D}}(d)^{\frac{\sigma-1}{\sigma}}\, dd\right]^{\frac{\sigma}{\sigma-1}} \quad \text{s.t.} \quad \text{Equation 17} \tag{23}$$

The planner's objective above is written for the closed economy. In the open economy, we characterize the planner's choices taking the realized export statuses as given.[19]

Distortion of the Outsourcing Share. The first theoretical result is that downstream firms in the decentralized economy *over-integrate*, i.e., they outsource less than is socially optimal, due to markups charged in firm-to-firm transactions. Conditional on output quality q^d, supplier quality q^u, and in-house quality q^h, the CES structure for material implies that the demand for an outsourced input u relative to the in-house input is

$$\begin{aligned}\frac{x^u(d)}{x^h(d)} &= \frac{\phi(q^d, q^u)}{\eta\phi(q^d, q^h)} \cdot \left(\frac{p^{\mathscr{U}}(q^u)}{p^h(q^h)}\right)^{-\sigma} \\ &= \frac{\phi(q^d, q^u)}{\eta\phi(q^d, q^h)} \left(\mu \frac{z(\omega^h, q^h)}{z(\omega^u, q^u)}\right)^{-\sigma}\end{aligned} \tag{24}$$

where μ is the suppliers' markup. The planner ignores μ, leading to a higher degree of outsourcing, holding other things fixed. Taking quality choices as given, the markup in Equation 24 distorts the intensive margin of material procurement.

Distortion of Quality Choice. To further characterize the distortion on the choice of quality investment, we compare the optimal quality choices of firms in the decentralized economy with those of the planner. The first-order condition for output quality q^d in the decentralized economy is:

$$wf'(q^d) = \frac{1}{\sigma} r^{\mathscr{D}}(d) \left[\frac{\partial \Phi(q^d)}{\partial q^d} \cdot \frac{1}{\Phi(q^d)} - \frac{\partial p^{\mathscr{D}}}{\partial q^d} \cdot \frac{\sigma-1}{p^{\mathscr{D}}}\right] \tag{25}$$

19. We do not characterize the efficient export decision, which depends on how the planner values export revenue; the objective above values only the domestic consumption bundle. For the same reason, the decentralized first-order conditions for quality in Equation 19 are stated conditional on a realized export status: the ex-ante condition also includes the derivative of expected export profit with respect to q^d, which operates through both revenue and the export probability.

whereas the planner's first-order condition, holding the input bundle fixed, is:

$$wf'(q^d) = \frac{1}{\sigma} r^{\mathscr{D}}_{\text{soc}}(d) \left[\frac{\partial \Phi(q^d)}{\partial q^d} \cdot \frac{1}{\Phi(q^d)} + \frac{\partial z}{\partial q^d} \cdot \frac{\sigma - 1}{z} \right] \tag{26}$$

Here $r^{\mathscr{D}}_{\text{soc}}(d)$ is the revenue of variety d in the planner's allocation, in which varieties are priced at marginal cost, whereas $r^{\mathscr{D}}(d)$ is the revenue in the market equilibrium, in which prices carry markups. Due to the markups, the decentralized firm's revenue is compressed relative to the planner's counterpart, a classic quantity distortion. We refer to this compressed revenue as the *scale effect*, which results in *under-investment* in output quality in the decentralized economy.

Relatedly, the decentralized downstream firm's first-order condition for in-house material quality q^h is:

$$\frac{\alpha_m(\sigma-1)}{\sigma} r^{\mathscr{D}}(d) \cdot \frac{\eta \phi(q^d, q^h)(p^h)^{1-\sigma}}{P_m^{1-\sigma}} \cdot \left[\frac{1}{\sigma - 1} \frac{\partial \phi(q^d, q^h)}{\partial q^h} \frac{1}{\phi(q^d, q^h)} - \frac{\partial p^h}{\partial q^h} \frac{1}{p^h} \right] = wf'(q^h) \tag{27}$$

whereas the planner's first-order condition is:

$$\underbrace{\frac{\alpha_m(\sigma-1)}{\sigma} \cdot r^{\mathscr{D}}_{\text{soc}}(d)}_{\text{Social Scale}} \cdot \underbrace{\frac{\eta \phi(q^d, q^h)(p^h)^{1-\sigma}}{P_{m,\text{soc}}^{1-\sigma}}}_{\text{Social In-house Share}} \cdot \left[\frac{1}{\sigma - 1} \frac{\partial \phi(q^d, q^h)}{\partial q^h} \frac{1}{\phi(q^d, q^h)} - \frac{\partial p^h}{\partial q^h} \frac{1}{p^h} \right] = wf'(q^h)$$

The in-house quality choice faces two countervailing distortions. On the one hand, the *scale effect* leads to under-investment: compressed revenue ($r^{\mathscr{D}} < r^{\mathscr{D}}_{\text{soc}}$) scales down the private return to any quality investment. On the other hand, markups make external inputs more expensive than in the planner's allocation ($P_m > P_{m,\text{soc}}$), so the in-house variety — priced at marginal cost in both worlds — is relatively *cheaper* in the decentralized economy. Its expenditure share within materials is correspondingly higher than the social in-house share, which raises the marginal return to in-house quality and induces over-investment. When this latter effect dominates, markups along supply chains lead to over-investment in in-house quality and thus over-integration relative to the efficient benchmark.

The potential over-investment in in-house quality remains even if we consider a different market structure from the monopolistically competitive one in firm-to-firm trade in our baseline. All else equal, under any type of rent-sharing with the supplier over ex post profits, the downstream firm pays more than marginal cost for external inputs, so the in-house variety remains relatively cheap and the distortion on its ex ante investment survives. This will play out when we move toward the policy counterfactuals discussed in Section 6.3.

4.4 Exposure to Suppliers' Markups

Section 4.3 asks how suppliers' markups change firms' choices. A second question is how much each firm's input bundle carries markups due to heterogeneous degrees of specialization. Substituting the downstream price Equation 9 into the sectoral price index Equation 11 gives:

$$p_s^{1-\sigma} = \int_{\mathscr{D}} \Phi(q^d) \left[\mu \cdot \frac{w^{\alpha_\ell} P_m(q^d, q^h)^{\alpha_m}}{z(\omega^d, q^d)} \right]^{1-\sigma} dd \tag{28}$$

To isolate the role of upstream markups, define the *shadow material price index* — the price index that would prevail if upstream inputs were priced at marginal cost rather than at a markup:

$$\tilde{P}_m^{1-\sigma} \equiv \int_{\mathscr{U}} \phi(q^d, q^u) \left[\frac{w}{z(\omega^u, q^u)} \right]^{1-\sigma} du \; + \; \eta\, \phi(q^d, q^h) \left[\frac{w}{z(\omega^h, q^h)} \right]^{1-\sigma} \tag{29}$$

The equilibrium material price index P_m departs from $\tilde{P}_m$ because external upstream inputs carry the constant markup μ, while the in-house variety is priced at marginal cost:

$$P_m^{1-\sigma} = \mu^{1-\sigma} \underbrace{\int_{\mathscr{U}} \phi(q^d, q^u) \left[\frac{w}{z(\omega^u, q^u)} \right]^{1-\sigma} du}_{\text{external (marked up)}} + \underbrace{\eta\, \phi(q^d, q^h) \left[\frac{w}{z(\omega^h, q^h)} \right]^{1-\sigma}}_{\text{in-house (no markup)}} \tag{30}$$

Let s^h denote the in-house share of material spending for a downstream firm d:

$$s^h = \frac{\eta\, \phi(q^d, q^h) \left[w / z(\omega^h, q^h) \right]^{1-\sigma}}{P_m^{1-\sigma}}$$

The equilibrium material price index can then be written as

$$P_m^{1-\sigma} = \tilde{P}_m^{1-\sigma} \cdot \underbrace{\left[\mu^{\sigma-1}(1 - s^h) + s^h \right]^{-1}}_{\equiv\, \Lambda(s^h)^{1-\sigma}} \tag{31}$$

where $\Lambda(s^h) \in [1, \mu]$ is the *double-marginalization factor*: the ratio of the equilibrium material price index to the shadow one, $P_m = \Lambda(s^h)\, \tilde{P}_m$. Under complete in-house production ($s^h = 1$), $\Lambda = 1$, meaning the upstream markup does not distort material costs. Conversely, under complete outsourcing ($s^h = 0$), $\Lambda = \mu$, and the full upstream markup passes through to the material price index. Between these extremes, Λ increases as expenditure shifts toward the marked-up external channel. Consequently, the burden of double marginalization is an endogenous outcome of the firm's make-or-buy decision. Holding its choices fixed, $\Lambda(s^h)^{\alpha_m}$ is the

firm's marginal cost relative to the case with every input priced at marginal cost. The markup is the same for every supplier; the exposure to it is not.

Substituting back and taking logs yields a three-way decomposition of the sectoral price index:

$$\ln p_s = \underbrace{\ln \mu}_{\substack{\text{downstream} \\ \text{markup}}} + \underbrace{\alpha_m \ln \bar{\Lambda}(\mathbf{s}^h)}_{\substack{\text{double} \\ \text{marginalization}}} + \underbrace{\ln \tilde{p}_s(q^d)}_{\substack{\text{quality-adjusted} \\ \text{marginal cost} \\ \text{(shadow price)}}} \tag{32}$$

where $\tilde{p}_s$ is the shadow sectoral price index constructed from the shadow downstream marginal costs $\tilde{p}^d = w^{\alpha_\ell} \tilde{P}_m^{\alpha_m} / z(\omega^d, q^d)$, and $\bar{\Lambda}(\mathbf{s}^h)$ aggregates the firm-level factors $\Lambda(s^h(d))$ across downstream firms, where $\mathbf{s}^h \equiv \{s^h(d)\}_{d \in \mathscr{D}}$ collects their in-house shares:

$$\bar{\Lambda}(\mathbf{s}^h)^{\alpha_m(1-\sigma)} \equiv \frac{\int_{\mathscr{D}} \Phi(q^d)\, \Lambda(s^h(d))^{\alpha_m(1-\sigma)}\, (\tilde{p}^d)^{1-\sigma}\, dd}{\int_{\mathscr{D}} \Phi(q^d)\, (\tilde{p}^d)^{1-\sigma}\, dd}$$

which is a power mean of Λ weighted by firms' expenditure shares in the shadow sectoral price index. Under this definition the decomposition holds exactly, and the contribution of each part on the right-hand side is of central interest to our quantitative application after parameter estimation.

Based on Equation 32, any shift in downstream firms' quality distribution affects the sectoral price index through two opposing channels. The *quality effect* operates through the shadow price: a higher q^d increases the final-demand weight $\Phi(q^d)$, which lowers $\tilde{p}_s$. However, due to the nonlinear profile of $z(\omega^d, q^d)$, higher quality raises productivity only as long as the elasticity $\omega_1^d + 2\bar{\omega} \log q^d$ remains positive, and reduces it thereafter. Furthermore, higher output quality reduces the weights $\phi(q^d, \cdot)$ on any given input mix, which increases the shadow material price index. This net quality effect is welfare-improving only when the demand-side gains dominate these supply-side forces. Conversely, the *double-marginalization effect* is unambiguously welfare-reducing. Due to the log-supermodularity of ϕ, a higher q^d tilts sourcing toward high-quality external suppliers. This shift reduces the in-house expenditure share and increases the double-marginalization factor Λ, which then passes through to the sectoral price index p_s.

5 Estimation

We estimate the export-only model of Section 4 following the method of simulated moments, where the main targets include the empirical patterns documented in Section 3. With the model estimated on the micro-level data of Taiwanese advanced manufacturing firms, our goal is to quantify the gains from outsourcing, to verify the mechanisms demonstrated in Section 4, and

to discuss policy counterfactuals that address the distortions described in Section 4.3. The welfare implications then follow from the decomposition in Section 4.4 using the estimated parameters.

Calibrated Parameters and Normalizations. The model has two groups of parameters. The first group is calibrated externally: the elasticity of substitution $\sigma = 2.5$, close to the average of the estimates of Aw and Lee (2017) for these four Taiwanese industries (2.3, with a range of 1.8 to 3.2 across them). This elasticity implies a monopolistically competitive markup of $\mu = \sigma/(\sigma - 1) = 1.67$, or 67% over marginal cost, for suppliers and downstream firms alike. The estimates of Aw and Lee (2017) come from Taiwanese firm-product data for the same industries, where the average 4-digit concentration (a Herfindahl–Hirschman index of about 0.05) is consistent with the monopolistic competition our constant markup assumes. Strictly, these estimates concern the demand for a firm's output, whereas the same σ also sets the markup in firm-to-firm trade through the material aggregator in Equation 6. We carry it over because the products of these industries are sold mostly to other producers, so the elasticity is largely one among intermediate varieties already, and because we observe no domestic transaction prices with which to identify a separate elasticity for the material bundle. The remaining calibrated parameters are the factor elasticities $\alpha_\ell = 0.3$ and $\alpha_m = 0.7$ calibrated to the aggregate statistics for these manufacturing sectors; and two normalizations, the wage $w = 1$ (the numeraire) and the labor endowment $\bar{L} = 1$. Aggregate income Y is determined in equilibrium by market clearing conditions.

Parameterization. The second group of thirteen parameters is estimated via the method of simulated moments (MSM). Here we go through the parameterization and the economic meaning of each parameter.

Every upstream firm draws ω_1^u from $\mathcal{LN}(\mu_\omega, \sigma_\omega^2)$, while each downstream firm draws its core and in-house upstream abilities (ω_1^d, ω_1^h) from a bivariate log-normal distribution with common mean μ_ω and dispersion σ_ω and correlation ϱ_ω. The parameters μ_ω, σ_ω, and ϱ_ω thus refer to the ability component ω_1. In addition, each production node draws the Hicks-neutral shifter $\omega_0 \overset{i.i.d.}{\sim} \mathcal{N}(0, \sigma_{\omega_0}^2)$ of Section 4.1. The productivity curvature $\bar{\omega}$ in Section 4.1 governs how strongly quality choice amplifies firm-level productivity differences.

The quality cost function is $f(q) = Fq^\gamma$, where F is the scale parameter and γ governs the curvature of the cost function with respect to quality. The quality complementarity parameter ν in $\phi(q^d, q^u)$ controls the strength of assortative matching between downstream and upstream quality. The parameter η is the weight on the in-house variety in the material aggregator. The foreign demand function is parameterized as $D_F(q^d) = b_1(q^d)^{b_2}$, where b_1 is a level shifter

and b_2 governs quality sensitivity. Finally, the export fixed cost is assumed to be lognormally distributed, $f_E \overset{i.i.d.}{\sim} \mathcal{LN}(\mu_E, \sigma_E^2)$. A higher μ_E will lower the share of exporting firms, while a higher σ_E will increase the dispersion of export status and export intensity among firms of similar quality.

Mapping Vertical Specialization: from Model to Data. In our estimation, we target a range of moments by specialization quintiles. To do this, we first need to reconcile vertical specialization in the model with its counterpart in the data.

In the model, sourcing distance takes only two values, zero steps for in-house production and one step for the upstream sector, so a firm's degree of specialization is inversely related to its in-house share of material spending, s^h. In the data, the specialization measure aggregates upstreamness gaps across many supplier industries. We bridge the two with the material-to-cost-of-goods-sold ratio that is observable in the data. In the model, it is also a monotone function of the in-house share:

$$\frac{\text{Material}}{\text{COGS}} = \alpha_m (1 - s^h)$$

To see this, note that under CRS ($\alpha_\ell + \alpha_m = 1$), cost-of-goods-sold (COGS) equals external purchased materials plus all labor, including downstream production labor and in-house upstream labor. Dividing external material expenditure by COGS yields exactly $\alpha_m \times (1 - s^h)$. As firms become more integrated (larger s^h), purchased materials replace a smaller share of COGS, so the ratio falls. Conversely, as firms become more specialized (smaller s^h), the ratio rises toward its upper bound α_m, which is the material intensity.

In the data, this mapping holds: specialized firms have a larger material share of COGS. Figure 7 plots the material share of COGS against $\text{Specialization}_{it}$. The more specialized a firm, the larger the share of its costs that goes to purchased materials. A one-standard-deviation increase in specialization is associated with a 5.4 log-point lower ratio of COGS to purchased materials, that is, a correspondingly larger material share of costs. This is also why specialized firms are more exposed to markups. In-house production carries no markup, while purchased materials do (Garetto 2013; Dhyne, Kikkawa, and Magerman 2022). Specialized firms therefore pay a markup on a larger share of their costs.

We therefore sort firms into quintiles of $\text{Specialization}_{it}$ in the data and of the in-house share s^h in the model, which is equivalent to sorting by the Material/COGS ratio. We include the material/COGS ratio across these quintiles as one set of targeted moments as well.

Moments and Identification. Table 5 summarizes the moment blocks alongside the parameters each mainly targets. We target 31 moments, computed on the 2017 cross-section of the analysis sample, with quintile moments demeaned by 6-digit industry. Five moment blocks are

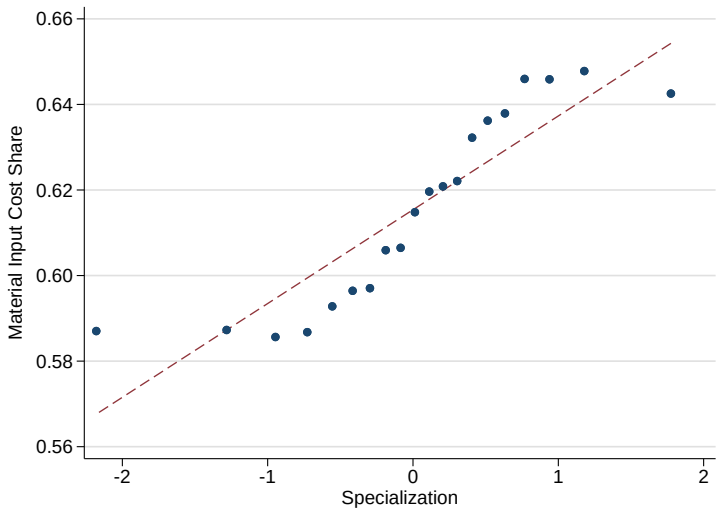


Figure 7: Material Share of Costs and Vertical Specialization

Note: Binscatter of the material share of COGS on Specialization$_{it}$, with both axes residualized on SIC6×year fixed effects (no firm controls in the figure). The coefficient reported in the text comes from regressing log COGS-to-materials on Specialization$_{it}$, controlling for log sales, log employment, firm age, and SIC6-by-year fixed effects, with standard errors clustered at the firm level: -0.054 (s.e. 0.006) on Specialization$_{it}$, $N = 94{,}854$.

computed as the average within each quintile of vertical specialization (quintile 1 the most specialized): the Material/COGS ratio, own log(1 + RDI), the transaction-share weighted supplier RDI, the share of exporting firms, and export intensity (exports over sales, with non-exporters at zero). The supplier-RDI profile is divided by its own first-quintile value in both model and data, so its first quintile equals one by construction and carries no weight; only the shape of the sorting profile is targeted.

The quintile moments pin down gradients across the specialization distribution but not levels, as they are demeaned within each 6-digit industry. We therefore add five aggregate levels: the Material/COGS ratio, the exporter share, mean export intensity among exporters, mean log(1 + RDI), and the pooled standard deviation of log sales, plus the OLS slope of log sales on the Material/COGS ratio, which disciplines the scale-specialization relationship.

In MSM every moment depends in principle on every parameter, and all 13 parameters are estimated jointly. That said, some moments are more important in identifying certain parameters, and the economic reasoning behind the mapping is as follows. The pooled standard deviation of log sales disciplines the dispersion of capability, σ_ω, together with the dispersion of the Hicks-neutral shifter, σ_{ω_0}: capability and quality choice alone generate too little sales dispersion, and the shifter supplies the remainder without altering the quality gradients. The OLS slope of log sales on the Material/COGS ratio disciplines how scale varies with special-

Table 5: Moments and Their Main Target Parameters

	Number of moments	Main target for
By quintile of vertical specialization (demeaned):		
Material/COGS ratio	5	ϱ_ω
log(1 + R&D/Sales)	5	ϱ_ω, γ
Supplier RDI (transaction-share weighted)	5	ν
Share of firms exporting	5	σ_E
Mean export intensity (export/sales)	5	b_2, $\bar{\omega}$
OLS slope of log sales on Material/COGS	1	μ_ω
Aggregate levels: Material/COGS; exporter share; export intensity (exporters); log(1 + R&D/Sales); S.D. of log sales	5	η, μ_E, b_1, F, σ_ω, σ_{ω_0}
Total	31	

Note: Quintile 1 is the most specialized. All 13 parameters are estimated jointly; the last column lists the parameters each block of moments is designed to discipline.

ization and is the main target for the location of capability, μ_ω. The productivity curvature $\bar{\omega}$ governs how much harder each unit is to produce at higher quality, which limits how much a high-quality firm sells; its main target is export intensity across quintiles, since foreign demand rises steeply with quality and export sales therefore respond most to it. The R&D intensity moments discipline the quality-cost function: the aggregate level of log(1 + RDI) disciplines the scale F, while the R&D gradient across quintiles disciplines the curvature γ, because more convex costs concentrate R&D spending among the highest-quality firms. The parameter γ is thus the curvature of the fixed cost of quality, whereas $\bar{\omega}$ is that of its variable cost. The share of firms exporting across quintiles pins down the dispersion of the export fixed cost, σ_E, and the aggregate exporter share its location, μ_E. Export intensity across specialization quintiles informs how steeply foreign demand rises with quality, b_2, and its aggregate level the shifter b_1.

The three parameters governing the make-or-buy margin — ϱ_ω, ν, and η — operate on distinct moments even though all three shape sourcing decisions. The correlation ϱ_ω is the main target of the gradients of the Material/COGS ratio and of R&D intensity across specialization quintiles: as the two capabilities become less positively, or even negatively, correlated, firms with strong core capability are more likely to draw weak in-house upstream capability, so they source externally and invest in output quality. The gradients therefore steepen as ϱ_ω falls. The complementarity parameter ν is identified by the *supplier-RDI gradient*, the direct empirical counterpart of Fact 2, since ν governs how strongly firms with high q^d demand high-q^u suppliers through ϕ. With $\nu = 0$, every buyer places the same weights on suppliers, so the

supplier-RDI profile is flat by construction; any gradient in it requires $\nu > 0$. The sorting profile is therefore what separates ν from ϱ_ω, since both steepen the Material/COGS gradient. The in-house weight η is identified by the aggregate *level* of the Material/COGS ratio: ϱ_ω redistributes in-house intensity across firms but leaves the sample average largely unchanged, so η is the parameter that scales the average sourcing pattern to match the data.

Procedure. Given a candidate parameter vector, we simulate 10,000 firms in each sector, discretize the quality choice on a grid of 46 evenly spaced points spanning $[0, 10]$, and solve the equilibrium by a nested fixed-point algorithm: an inner loop solves the upstream price and demand schedules and the sectoral price index p_s at $w = 1$, and an outer loop adjusts aggregate income Y until labor demand equals the endowment $\bar{L} = 1$. The estimator minimizes the GMM quadratic form $g'Wg$, where g stacks the model-minus-data moment deviations and W is the inverse of the bootstrapped variance–covariance matrix of the data moments, the efficient weighting matrix, using the Nelder–Mead simplex. We describe the algorithm and the weighting matrix in Appendix B.

Parameter Estimates and Fit. Table 6 reports the estimates and their standard errors. Several features stand out. First, the capability correlation is negative ($\varrho_\omega = -0.41$): firms with strong core quality capability tend to draw weak in-house upstream capability. The sign deserves a caveat. At first glance it may read as an absolute disadvantage of high-quality firms at making their inputs, but the mechanism in our model requires only comparative advantage: any $\varrho_\omega < 1$ leaves some firms better at upgrading their output than their inputs (Section 4.1). The estimate is negative because quality brings scale: sales rise with output quality, so high-quality firms are large and can more easily cover the fixed cost of in-house quality. On scale alone, they would integrate more, not less. Matching the data therefore requires comparative advantage strong enough to offset this scale advantage, and that is what the estimation delivers. A negative correlation can also be rationalized by a span-of-control model in which attention diverted from the core business carries strongly diminishing returns, particularly for the firms most capable in their main product.

Second, the complementarity parameter is $\nu = 0.43$: to keep an input equally useful, its quality must rise by 0.43 units for each unit increase in the buyer's output quality, as explained in Section 4.1. To put this in terms of input sourcing, in the estimated economy a buyer in the top decile of output quality places a 38% higher weight on a top-decile supplier than on a bottom-decile one, against 13% for a buyer in the bottom decile. At equal prices, these are the relative quantities the two buyers purchase. Third, the productivity curvature is negative ($\bar{\omega} = -0.29$), generating diminishing returns to quality in productivity, while the quality cost

curvature $\gamma = 7.8$ makes quality investment strongly convex.[20] Fourth, the in-house weight is below one ($\eta = 0.09$): in-house production supplies the variety-equivalent of roughly a tenth of the external market, in line with the measure reading of η introduced with the material aggregator. Beyond scope, η captures any residual costs of in-house input production, such as setup costs, managerial scope, and coordination overhead. The magnitude of this parameter is disciplined by a specific empirical moment: the aggregate mean Material/COGS ratio.

Table 6: Parameter Estimates

Description	Parameter	Estimate	SE
Mean of $\log \omega_1$	μ_ω	0.679	0.031
S.D. of $\log \omega_1$	σ_ω	0.493	0.018
Quality curvature	$\bar{\omega}$	−0.292	0.019
S.D. of Hicks-neutral shifter	σ_{ω_0}	0.336	0.018
Capability correlation	ϱ_ω	−0.411	0.020
Quality-cost scale (10^{-9})	F	2.41	0.625
Quality-cost curvature	γ	7.780	0.589
Quality complementarity	ν	0.429	0.046
In-house weight	η	0.092	0.004
Foreign demand shifter (10^{-6})	b_1	1.84	0.226
Foreign demand curvature	b_2	3.067	0.262
Mean of $\log f_E$	μ_E	−11.737	0.116
S.D. of $\log f_E$	σ_E	2.218	0.171

Note: Method of simulated moments estimates; 31 moments, 13 parameters. Standard errors are computed from the efficient-MSM formula $(G'\hat{\Omega}^{-1}G)^{-1}$, where G is the Jacobian of the model moments at the estimates (central differences) and $\hat{\Omega}$ is the estimated moment covariance matrix used as the weighting matrix.

Table 7 reports the full set of data and model moments at the estimated parameters. The model reproduces the gradients that carry the paper's mechanism (the Material/COGS decline across quintiles, the R&D intensity gradient, the supplier-RDI sorting gradient, and the export participation gradient) as well as the aggregate Material/COGS ratio, the aggregate exporter share, and the aggregate export intensity among exporters.

20. The convexity is anecdotally evident in the sectors we study, and arguably why firms in these sectors adopted a business model of vertical specialization in recent decades, especially for electronics sectors after Morris Chang founded TSMC (Thurk 2022).

Table 7: Model Fit: MSM Moments by Specialization Quintile

Moment		Q1	Q2	Q3	Q4	Q5
Material/COGS	Data	0.021	0.028	0.008	−0.016	−0.041
	Model	0.052	0.037	0.020	−0.008	−0.101
Share exporting	Data	0.048	0.080	0.039	−0.024	−0.143
	Model	0.082	0.030	−0.001	−0.033	−0.078
Export intensity	Data	0.037	0.054	0.015	−0.025	−0.081
	Model	0.056	0.011	−0.009	−0.024	−0.035
log(1 + RDI)	Data	0.0148	0.0017	−0.0021	−0.0065	−0.0080
	Model	0.0159	0.0028	−0.0020	−0.0060	−0.0106
Supplier RDI (profile, Q1 = 1)†	Data	1.000	−0.096	−0.290	−0.367	−0.249
	Model	1.000	0.178	−0.151	−0.398	−0.628

	Data	Model
OLS slope: log sales on Material/COGS	1.123	0.936
Aggregate Material/COGS	0.597	0.630
Aggregate exporter share	0.591	0.608
Aggregate export intensity (exporters)	0.251	0.286
Aggregate log(1 + RDI)	0.026	0.040
Aggregate S.D. of log sales	1.642	1.182

Note: Quintiles are sorted by vertical specialization in the data (Q1 = most specialized). Quintile moments are demeaned by 6-digit industry in the data and by the simulation mean in the model. †The supplier-RDI profile is divided by its own Q1 value in both model and data, so Q1 = 1 by construction and carries no weight; only the shape of the profile is targeted. Aggregate levels are pooled means; the pooled standard deviation of log sales is targeted as an aggregate.

6 Quantitative Exercise

We use the estimated model for three exercises. Section 6.1 separates the roles of comparative advantage and complementarity in the link between output quality and outsourcing. Section 6.2 measures how exposure to suppliers' markups varies with output quality. Section 6.3 evaluates two policies toward these markups: removing them, and a credit on input purchases.

6.1 Comparative Advantage Generates Specialization; Complementarity Amplifies It

Two estimated parameters govern the link between output quality and outsourcing: the correlation ϱ_ω between a firm's core and in-house capabilities, and the complementarity parameter ν. To separate their roles, we re-solve the equilibrium twice, once with $\nu = 0$ and once with $\varrho_\omega = 0$, holding the other parameters at their estimates and the exchange rate at $e = 1$.

Figure 8 plots each simulated firm's material share of costs, which in our theory reflects the degree of vertical specialization, against its chosen output quality q^d, the model counterpart of Figure 7, with an OLS line for each economy.

In the baseline, higher-quality firms buy more of their inputs: the slope is positive, with a 95% confidence interval of $[0.010, 0.013]$. Comparative advantage generates this link. With $\varrho_\omega = 0$, the slope turns negative, $[-0.012, -0.009]$: if the two capabilities were uncorrelated, high-quality firms could afford in-house quality spending with their scale, and would outsource less. Complementarity amplifies the link. With $\nu = 0$, the slope flattens to $[0.001, 0.004]$: output quality no longer raises the demand for input quality, so high-quality firms outsource only slightly more and no longer sort toward higher-quality suppliers.

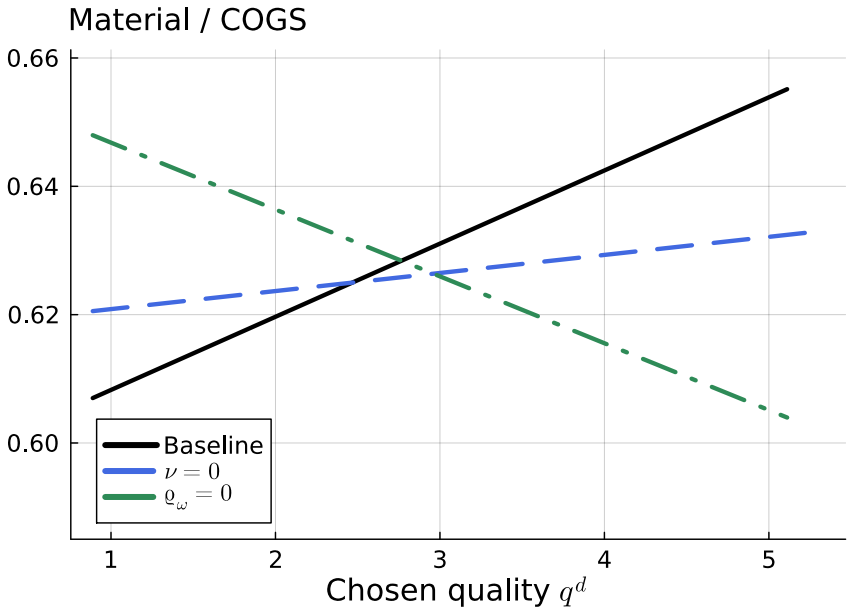


Figure 8: Output Quality and Outsourcing: Comparative Advantage and Complementarity

Note: Material/COGS against chosen output quality q^d in the simulated cross-section of the estimated model (baseline), the model with $\nu = 0$, and the model with $\varrho_\omega = 0$. Lines are OLS fits over the 2nd to 98th percentiles of q^d, on the firms active in all three economies. 95% confidence intervals of the slopes: baseline $[0.010, 0.013]$; $\nu = 0$: $[0.001, 0.004]$; $\varrho_\omega = 0$: $[-0.012, -0.009]$.

6.2 Exposure to Suppliers' Markups Rises with Output Quality

Because high-quality firms outsource more, a larger share of their costs carries suppliers' markups. We measure this exposure by $\Lambda(s^h)^{\alpha_m}$ from Section 4.4: a firm's marginal cost relative to the same firm with every input priced at marginal cost, holding its choices fixed. Figure 9 plots it against output quality. Across deciles of q^d, the in-house share falls from 14%

to 8%, so suppliers' markups put the bottom decile's costs 38% above marginal cost and the top decile's 40%.[21] Although the markup is constant, exposure to it varies with quality because quality is correlated with specialization; the resulting dispersion in exposure distorts relative prices across firms, a source of misallocation in these advanced manufacturing sectors.

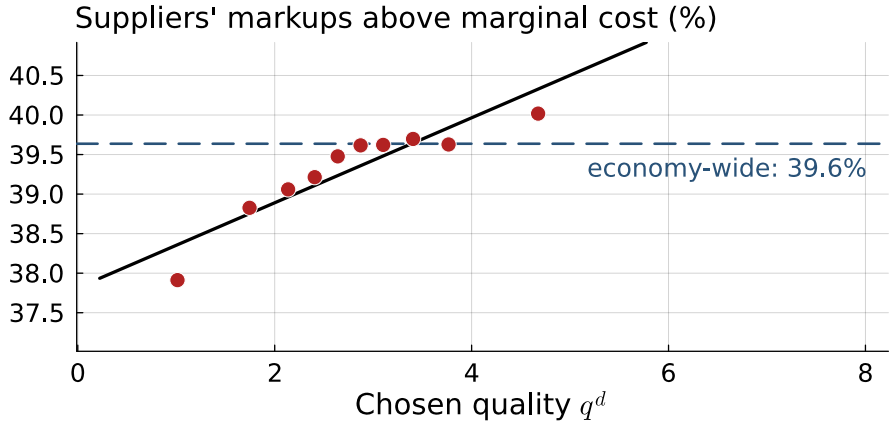


Figure 9: Exposure to Suppliers' Markups and Output Quality

Note: Suppliers' markups above marginal cost, $\Lambda(s_i^h)^{\alpha_m} - 1$, against chosen output quality q^d in the simulated baseline. Dots: means within deciles of q^d (about 1,000 firms each), across which the in-house share falls from 14% to 8%. Solid line: OLS fit across firms. Dashed line: the economy-wide value, $\bar{\Lambda}^{\alpha_m} - 1 = 39.6\%$, which weights firms by their shares in the sectoral price index.

Aggregating over firms gives the decomposition of the sectoral price index in Equation 32, reported in column (1) of Table 8. The downstream markup contributes $\ln \mu = 0.511$; this term is constant across counterfactuals. Suppliers' markups contribute $\alpha_m \ln \bar{\Lambda} = 0.334$: they put the economy's costs $e^{0.334} - 1 \approx 40\%$ above marginal cost, holding quality choices fixed. The quality component contributes $\ln \tilde{p}_s = -1.391$: quality investment along the chain lowers the quality-adjusted marginal cost of the final good by about 75%. The next subsection asks whether policy can lower the cost of suppliers' markups without giving up the quality they help finance.

6.3 Policy and the Dual Role of Suppliers' Markups

Now that we have established the relationship between quality and specialization, and the exposure to markups that comes with it, a natural question is: what policy can reduce the cost of

21. For reference, with $\sigma = 2.5$ a supplier's markup is $\mu - 1 = 67\%$; with materials a share $\alpha_m = 0.7$ of costs, a firm that bought every input would have costs $\mu^{\alpha_m} - 1 = 43\%$ above marginal cost, and a firm that made every input none. The estimated in-house shares place firms at 38–40%.

using the market without sacrificing the gains from quality investment, especially when complementarity between input and output quality is as pervasive as in advanced manufacturing?

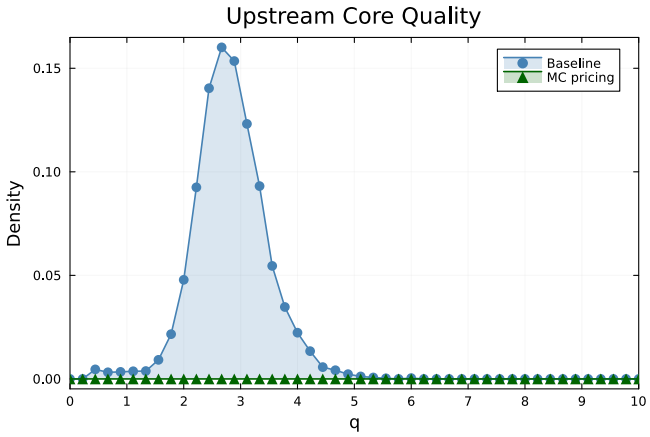
Recent work on industrial policy prescribes that, because distortions compound through input-output linkages, policy should remove them upstream (Liu 2019). In quality-critical supply chains, however, suppliers' markups also play a second role: they reward upstream firms for investing in quality. In the model, this role appears in the supplier's first-order condition in Section 4.1: the operating margin σ^{-1} is what turns quality-driven demand into profit. Once markups are not pure distortions, any intervention that affects the division of surplus ex post also shapes investment ex ante (Grossman and Hart 1986; Acemoglu et al. 2010).

We run two exercises, each re-solving the equilibrium at the estimated parameters. The first removes the markup, pricing external inputs at marginal cost. The second subsidizes the purchase, reimbursing part of every external input bill and financing it by a lump-sum charge on households. Table 8 summarizes both.

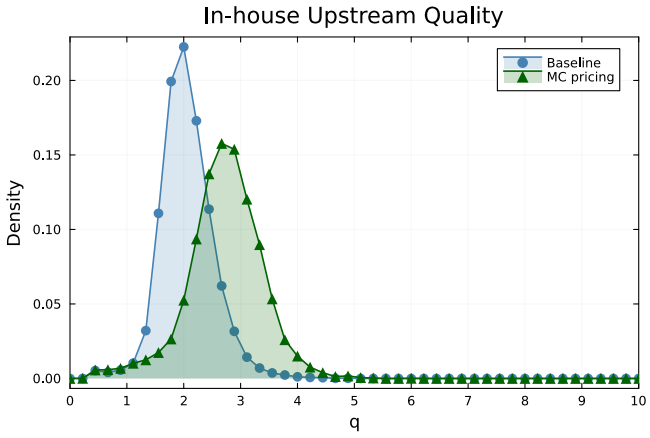
Quantifying the Gains from Specialization. In our first counterfactual exercise, where upstream markups are removed, we also have the opportunity to quantify the gains from outsourcing as downstream firms are forced to perform upstream tasks entirely in-house. With external inputs priced at marginal cost, upstream firms keep no operating margin to cover the fixed cost of quality, so they do not invest and stay out of the market, and downstream firms make all their inputs in-house ($s^h = 1$). This also sets the double-marginalization term in the sectoral price index to zero.

Figure 10 illustrates the consequences for product quality. In the downstream sector, in-house input quality rises because firms are now forced to produce all inputs internally. This directly diverts investment away from output quality. Indirectly, because high-quality downstream firms tend to lack a comparative advantage in in-house quality upgrading, the return on output-quality investment declines due to the presence of complementarity. A firm's input quality becomes constrained by its internal capabilities rather than by superior suppliers. Consequently, average output quality falls, and the upper tail of the quality distribution collapses. Although eliminating supplier markups saves the 0.334 log points that these markups mechanically add to the sectoral price index, the quality component rises by 1.43 log points. In sum, the sectoral price index rises by 198% and real consumption, $C = Y/p_s$, falls by 40% (column (2) of Table 8).[22] Our baseline economy, while subject to double marginalization, still realizes a large gain from specialization.

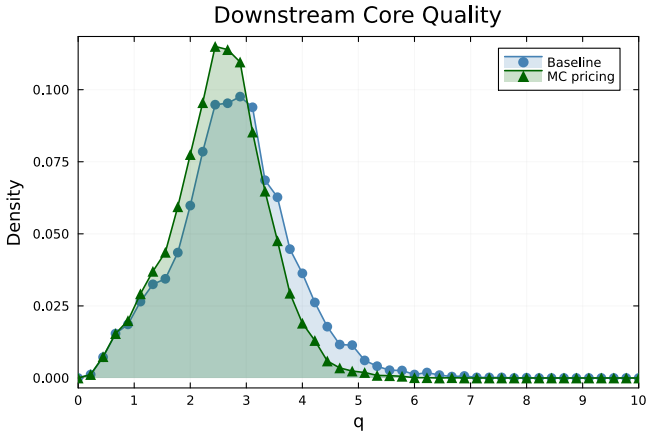
22. This is computed at the baseline exchange rate, $e = 1$, at which real exports fall by 97%.



(a) Supplier quality q^u



(b) In-house quality q^h



(c) Output quality q^d

Figure 10: Quality When External Inputs Are Priced at Marginal Cost

Note: Distributions of chosen quality in the estimated baseline and in the counterfactual in which external inputs are priced at marginal cost ("MC pricing"). Upstream firms exit in the counterfactual, so panel (a) has no counterfactual mass. The quality space is discretized on $[0, 10]$.

A Credit on Input Purchases. The second counterfactual exercise speaks to recent industrial policies, in Taiwan and elsewhere, that aim to foster the development of advanced manufacturing supply chains. Since 2023, Taiwan's Statute for Industrial Innovation has allowed firms to deduct 5% of their spending on advanced-process equipment from corporate income tax (Article 10-2).[23] Our counterfactual applies a credit of this kind to firm-to-firm input purchases: the government reimburses a fraction t of every external input bill and finances it by a lump-sum charge on households. The buying downstream firm pays 5% less, while the supplier still receives the full markup and keeps its margin. In-house production receives nothing. We report two cases: $t = 5\%$, Taiwan's rate for equipment procurement, and the full offset $t^* = 1 - 1/\mu = 0.40$, at which buyers pay marginal cost. In both, the exchange rate adjusts to hold real exports at their baseline level.

We summarize the decomposition of the sectoral price index and the counterfactual real consumption in column (3) of Table 8. A 5% credit raises real consumption by 2.7%, net of the lump-sum charge of 2% of nominal income that finances it. It lowers the contribution of suppliers' markups to the sectoral price index from 0.334 to 0.302 log points, and the in-house share from 8.4% to 7.7%. The gain comes mostly from lowering the markup on all purchased inputs. The full offset, summarized in column (4) of Table 8, raises real consumption by 18.3%. Buyers then face no markup, the in-house share falls to 3.5%, and the exchange rate appreciates to $e = 0.75$ to hold real exports at their baseline level. In-house quality falls, supplier quality rises at the top of the distribution, and output quality is slightly lower on average (see the plots in Section C.2). While the gain is large, the full offset requires a transfer of 27% of nominal income, which may be politically infeasible.

Together, the two exercises show that suppliers' markups are not a pure distortion. They discourage outsourcing, and they raise costs most for the high-quality firms that still outsource the most; but they also fund the supplier quality on which downstream firms depend. Specialization, despite markup distortions, is worth 40% of real consumption. A credit on input purchases lowers what buyers pay while suppliers keep their margin, so it raises welfare; forcefully removing the markups lowers it.

7 Conclusion

This paper studies how quality is produced across firm boundaries when input and output quality are strongly complementary, as in advanced manufacturing, and quantifies the gains and costs from outsourcing. In Taiwan, we show that the most vertically specialized firms produce the highest-quality output and buy from the suppliers that invest the most in quality.

23. The text of the Statute, in Chinese, is available from the Ministry of Justice's Laws and Regulations Database.

Table 8: Welfare: Removing Suppliers' Markups vs. a Credit on Input Purchases

	Baseline	No markup	Credit $t = 5\%$	Credit $t^* = 40\%$
	(1)	(2)	(3)	(4)
$\ln \mu$ (downstream markup)	0.511	0.511	0.511	0.511
$\alpha_m \ln \bar{\Lambda}$ (suppliers' markups)	0.334	0.000	0.302	0.000
$\ln \tilde{p}_s$ (quality-adjusted marginal cost)	−1.391	0.035	−1.392	−1.393
$\ln p_s$	−0.546	0.545	−0.580	−0.882
Final-good price index p_s	0.579	1.725	0.560	0.414
Average in-house share s^h	0.084	1.000	0.077	0.035
Upstream firms active	yes	none	yes	yes
Exchange rate e	1.000	1.000	0.972	0.748
Real exports X/p_s	1.453	0.047	1.458	1.458
Credit outlay / income	—	—	2%	27%
Real consumption $C = (w\bar{L} + \Pi - X)/p_s$	1.355	0.816	1.391	1.602
Δ vs. baseline	—	−39.8%	+2.7%	+18.3%

Note: Each column re-solves the equilibrium at the estimated parameters. Column (2) prices external inputs at marginal cost. Columns (3) and (4) reimburse a fraction t of every external input bill, financed by a lump-sum charge on households. Columns (3) and (4) let the exchange rate adjust to hold real exports at their baseline level; column (2) holds $e = 1$, because no exchange rate sustains baseline real exports once quality falls dramatically.

We attribute this pattern to comparative advantage and complementarity between input and output quality. In the estimated model, high-quality firms outsource the most, so more of their costs carry suppliers' markups than is the case for more integrated firms.

These markups discourage outsourcing, but they also fund supplier quality investment. Pricing external inputs at marginal cost lowers welfare by forcing every firm to integrate: suppliers exit, firms make all their inputs in-house, output quality falls, and real consumption falls by 40%. A credit on input purchases instead lowers what buyers pay while suppliers keep their margin. Applied at the 5% rate Taiwan grants on advanced-process equipment and financed by a lump-sum charge on households, it raises real consumption by 2.7%.

Three potential extensions follow. First, the markup could be microfounded, as the outcome of hold-up problems or of imperfect competition among suppliers (Grossman and Hart 1986; Dhyne, Kikkawa, and Magerman 2022). Second, trade policy could be studied in an economy whose upstream sector is still immature, where import substitution and export promotion shape how that sector develops and how firms' vertical boundaries are determined.

Third, a dynamic version could let demand for quality rise with income and let firms choose among technologies with different degrees of complementarity. We leave these to future work.

References

Acemoglu, Daron, Rachel Griffith, Philippe Aghion, and Fabrizio Zilibotti. 2010. "Vertical Integration and Technology: Theory and Evidence." *Journal of the European Economic Association* 8 (5): 989–1033.

Antràs, Pol. 2003. "Firms, Contracts, and Trade Structure." *The Quarterly Journal of Economics* 118 (4): 1375–1418.

Antràs, Pol, Davin Chor, Thibault Fally, and Russell Hillberry. 2012. "Measuring the Upstreamness of Production and Trade Flows." *American Economic Review* 102 (3): 412–416.

Arkolakis, Costas, Federico Huneeus, and Yuhei Miyauchi. 2023. *Production Network Formation, Trade, and Welfare.* NBER Working Paper 30954. National Bureau of Economic Research.

Atalay, Enghin, Ali Hortaçsu, Mary Jialin Li, and Chad Syverson. 2019. "How Wide Is the Firm Border?" *The Quarterly Journal of Economics* 134 (4): 1845–1882.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson. 2014. "Vertical Integration and Input Flows." *American Economic Review* 104 (4): 1120–1148.

Aw, Bee Yan, and Yi Lee. 2017. "Demand, Costs and Product Scope in the Export Market." *European Economic Review* 100:28–49.

Becker, Gary S., and Kevin M. Murphy. 1992. "The Division of Labor, Coordination Costs, and Knowledge." *The Quarterly Journal of Economics* 107 (4): 1137–1160.

Bernard, Andrew B., Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes. 2022. "The Origins of Firm Heterogeneity: A Production Network Approach." *Journal of Political Economy* 130 (7): 1765–1804.

Bernard, Andrew B., and Andreas Moxnes. 2018. "Networks and Trade." *Annual Review of Economics* 10:65–85.

Bijnens, Gert, Anders Humlum, and Stijn Vanormelingen. 2026. *Capital-Skill Complementarity Beyond Firm Boundaries.* NBER Working Paper 35857. National Bureau of Economic Research.

Boehm, Johannes, and Ezra Oberfield. 2020. "Misallocation in the Market for Inputs: Enforcement and the Organization of Production." *The Quarterly Journal of Economics* 135 (4): 2007–2058.

———. 2023. *Growth and the Fragmentation of Production.* Working Paper.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *The Review of Economic Studies* 89 (1): 181–213.

Broda, Christian, and David E. Weinstein. 2006. "Globalization and the Gains from Variety." *The Quarterly Journal of Economics* 121 (2): 541–585.

Cajal-Grossi, Julia, Rocco Macchiavello, and Guillermo Noguera. 2023. "Buyers' Sourcing Strategies and Suppliers' Markups in Bangladeshi Garments." *The Quarterly Journal of Economics* 138 (4): 2391–2450.

Chen, Liang-Chih. 2014. "Entrepreneurship, Technological Changes, and the Formation of a Subcontracting Production System: The Case of Taiwan's Machine Tool Industry." *International Journal of Economics and Business Research* 7 (2): 198–219.

Chen, Teddy Chia-Wei, Joonkyo Hong, Toshiaki Komatsu, Jin-Tan Liu, and Jen-Kuan Wang. 2026. *Taiwanese Domestic Supply Chains: 2013–2022.* Technical report. National Taiwan University.

Coase, Ronald H. 1937. "The Nature of the Firm." *Economica* 4 (16): 386–405.

Costinot, Arnaud. 2009. "An Elementary Theory of Comparative Advantage." *Econometrica* 77 (4): 1165–1192.

Cournot, Antoine Augustin. 1838. *Recherches sur les principes mathématiques de la théorie des richesses.* Paris: Hachette.

Demir, Banu, Ana Cecília Fieler, Daniel Yi Xu, and Kelly Kaili Yang. 2024. "O-Ring Production Networks." *Journal of Political Economy* 132 (1): 200–247.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, and Glenn Magerman. 2022. "Imperfect Competition in Firm-to-Firm Trade." *Journal of the European Economic Association* 20 (5): 1933–1970.

Dhyne, Emmanuel, Ayumu Ken Kikkawa, Magne Mogstad, and Felix Tintelnot. 2021. "Trade and Domestic Production Networks." *The Review of Economic Studies* 88 (2): 643–668.

Eaton, Jonathan, Marcela Eslava, David Jinkins, C. J. Krizan, and James R. Tybout. 2025. "A Search and Learning Model of Export Dynamics." *Journal of International Economics* 157:104155.

Eaton, Jonathan, David Jinkins, James R. Tybout, and Daniel Xu. 2022. *Two-Sided Search in International Markets.* NBER Working Paper 29684. National Bureau of Economic Research.

Eaton, Jonathan, Samuel S Kortum, and Francis Kramarz. 2022. *Firm-to-Firm Trade: Imports, exports, and the labor market.* Technical report. National Bureau of Economic Research.

Ekerdt, Lorenz K. F., and Kai-Jie Wu. 2024. *The Rise of Specialized Firms.* Working Paper 24-06. Center for Economic Studies, U.S. Census Bureau.

Fieler, Ana Cecília, Marcela Eslava, and Daniel Yi Xu. 2018. "Trade, Quality Upgrading, and Input Linkages: Theory and Evidence from Colombia." *American Economic Review* 108 (1): 109–146.

Ganapati, Sharat, and Colin J. Hottman. 2026. *Did Foreigners Pay America's Tariffs? Quantity Discounts, Scale Economies and Incomplete Pass-Through.* NBER Working Paper 34901. National Bureau of Economic Research.

Garetto, Stefania. 2013. "Input Sourcing and Multinational Production." *American Economic Journal: Macroeconomics* 5 (2): 118–151.

Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy* 94 (4): 691–719.

Hansman, Christopher, Jonas Hjort, Gianmarco León-Ciliotta, and Matthieu Teachout. 2020. "Vertical Integration, Supplier Behavior, and Quality Upgrading among Exporters." *Journal of Political Economy* 128 (9): 3570–3625.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang. 2014. "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." *American Economic Review* 104 (6): 1597–1629.

Huneeus, Federico. 2018. *Production Network Dynamics and the Propagation of Shocks.* Working Paper.

Johnson, Robert C., and Andreas Moxnes. 2023. "GVCs and Trade Elasticities with Multistage Production." *Journal of International Economics* 145:103796.

Khandelwal, Amit. 2010. "The Long and Short (of) Quality Ladders." *The Review of Economic Studies* 77 (4): 1450–1476.

Khandelwal, Amit K., Peter K. Schott, and Shang-Jin Wei. 2013. "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters." *American Economic Review* 103 (6): 2169–2195.

Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *The Quarterly Journal of Economics* 108 (3): 551–575.

Kugler, Maurice, and Eric Verhoogen. 2012. "Prices, Plant Size, and Product Quality." *The Review of Economic Studies* 79 (1): 307–339.

Lim, Kevin. 2018. *Endogenous Production Networks and the Business Cycle.* Working Paper.

Liu, Ernest. 2019. "Industrial Policies in Production Networks." *The Quarterly Journal of Economics* 134 (4): 1883–1948.

Macchiavello, Rocco, Josepa Miquel-Florensa, Nicolas de Roux, Eric Verhoogen, Mario Bernasconi, and Patrick W. Farrell. 2026. *Quality Upgrading in Global Supply Chains: Evidence from Colombian Coffee.* NBER Working Paper 34610. National Bureau of Economic Research.

Ministry of Economic Affairs. 2005. *2005 Industrial Statistical Survey Report.* Technical report. Ministry of Economic Affairs, Taiwan.

Miyauchi, Yuhei. 2024. "Matching and Agglomeration: Theory and Evidence From Japanese Firm-to-Firm Trade." *Econometrica* 92 (6): 1869–1905.

Riordan, Michael H., and Oliver E. Williamson. 1985. "Asset Specificity and Economic Organization." *International Journal of Industrial Organization* 3 (4): 365–378.

Schumpeter, Joseph A. 1934. *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle.* Translated by Redvers Opie. Cambridge, MA: Harvard University Press.

Spengler, Joseph J. 1950. "Vertical Integration and Antitrust Policy." *Journal of Political Economy* 58 (4): 347–352.

Sutton, John. 2001. *Technology and Market Structure: Theory and History*. Cambridge, MA: The MIT Press.

Thurk, Jeff. 2022. *Outsourcing, Firm Innovation, and Industry Dynamics in the Production of Semiconductors.* Working Paper.

Tsai, Yiu Hing Barron. 2025. *When Cutting Out the Middleman Backfires: Disintermediation, Wholesale Markups, and Misallocation.* Working Paper.

Viratyosin, Tasaneeya. 2026. *Innovation Spillovers Along the Supply Chain.* Working Paper.

Yi, Kei-Mu. 2003. "Can Vertical Specialization Explain the Growth of World Trade?" *Journal of Political Economy* 111 (1): 52–102.

A Additional Empirical Results

A.1 Export and Import Unit Prices

As a complement to the demand-residual quality measure in Fact 1, we examine unit prices in the customs data. For firm i, HS6 product h, destination country c, and year t, we estimate

$$\log\big(\text{USD/kg}\big)_{ihct} = \beta_1 \cdot \text{Specialization}_{it} + \beta_2 \cdot \log \text{Sales}_{ihct} + \mathbf{X}'_{it}\delta + \lambda_{hct} + \mu_{k(i)t} + \varepsilon_{ihct},$$

where $\log \text{Sales}_{ihct}$ is the firm's log export sales (or import purchases) in that product-destination-year market, $\mathbf{X}_{it}$ contains log sales, log employment, and firm age, λ_{hct} is an HS6×destination×year fixed effect, and $\mu_{k(i)t}$ is a SIC6-by-year fixed effect. In the first column of each panel of Table A1, we drop the market-share control $\log \text{Sales}_{ihct}$; in the second column we include it. The results indicate that specialized firms charge higher unit prices in export markets and incur higher unit costs for imports, even when accounting for their total sales or purchases in those markets.

Table A1: Vertical Specialization and Export and Import Pricing

	Exports		Imports	
	$\log$ (USD/kg)$_{ihct}$		$\log$ (USD/kg)$_{ihct}$	
Specialization$_{it}$	0.193	0.235	0.122	0.143
	(0.019)	(0.019)	(0.017)	(0.017)
$\log$ Sales$_{ihct}$		0.134		0.144
		(0.003)		(0.003)
HS6-Country-Year FE	✓	✓	✓	✓
SIC6-Year FE	✓	✓	✓	✓
Firm controls	✓	✓	✓	✓
Observations	1,850,323	1,850,323	1,048,573	1,048,573
R^2	0.425	0.455	0.490	0.518

Note: For each HS6 product h exported to/imported from country c in year t, the per kilogram price is higher for specialized firms, controlling for total values in that market. Specialization is the negative of vertical span, standardized within industry-year. $\log \text{Sales}_{ihct}$ is the firm's log export sales (left panel) or import purchases (right panel) in the product-destination-year market. Controls: log sales, firm age, log employment. Standard errors in parentheses are clustered at the firm level.

A.2 Supplier Sorting: Alternative Quality Measures

Table A2 repeats the sorting regressions of Fact 2 (with own quality entering additively, interacted with Specialization$_{it}$, and controls plus $\mu_{k(i)t}$ as before) for two alternative measures of own quality: the extensive margin of R&D (an indicator for positive R&D spending, with supplier quality measured by the share of purchases from R&D-performing suppliers) and log R&D intensity. The two headline patterns of the baseline table carry over to both specifications: firms sort positively with their suppliers on every margin, and more specialized firms source from higher-quality suppliers at every level of own quality. The interaction between own quality and specialization, by contrast, is positive in the extensive-margin specification but negative in the log specification.

Table A2: R&D Sorting: Extensive-Margin and Log Specifications

	Supplier $\mathbf{1}\{\text{RD} > 0\}_{it}$			Supplier $\log \text{RDI}_{it}$		
	Weighted	Avg	Cov	Weighted	Avg	Cov
Quality$_{it}$	0.070	0.021	0.050	0.093	0.046	0.047
	(0.004)	(0.001)	(0.004)	(0.011)	(0.005)	(0.009)
Specialization$_{it}$	0.014	0.006	0.009	0.135	0.102	0.033
	(0.002)	(0.000)	(0.002)	(0.042)	(0.019)	(0.036)
Quality$_{it}$ × Specialization$_{it}$	0.068	0.009	0.059	−0.028	−0.005	−0.023
	(0.006)	(0.001)	(0.006)	(0.013)	(0.005)	(0.011)
SIC6×Year FE	✓	✓	✓	✓	✓	✓
Controls	log Sales, Age, log Employment					
Observations	116,488			18,169		
R^2	0.236	0.285	0.175	0.142	0.190	0.082

Note: Quality $= \mathbf{1}\{\text{RD}_{it} > 0\}$ (left panel) and $\log \text{RDI}_{it}$ (right panel). In the left panel, supplier quality is the share of purchases from R&D-performing suppliers; in the right panel, it is supplier log R&D intensity. Weighted = transaction-share weighted supplier quality; Avg = unweighted mean supplier quality; Cov = Weighted − Avg, the covariance between transaction shares and supplier quality. Controls: log sales, firm age, log employment; SIC6×year fixed effects. Standard errors in parentheses are clustered at the firm level.

A.3 Correlates of Vertical Specialization

Table A3 reverses the regressions of Table 2: it regresses Specialization$_{it}$ on each quality-related characteristic, controlling for log sales, log employment, firm age, and log wage, with SIC6-by-year fixed effects. Section 3 discusses the results.

Table A3: Correlates of Vertical Specialization

	Dependent variable: Specialization$_{it}$					
	(1)	(2)	(3)	(4)	(5)	(6)
log Sales$_{it}$	0.057	0.128	0.049	0.059	0.054	0.017
	(0.007)	(0.016)	(0.007)	(0.008)	(0.007)	(0.022)
log Emp$_{it}$	−0.066	−0.189	−0.057	−0.077	−0.049	−0.034
	(0.008)	(0.020)	(0.008)	(0.010)	(0.008)	(0.025)
Age$_{it}$	−0.010	−0.008	−0.010	−0.009	−0.010	−0.010
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
log wage$_{it}$	0.137	0.283	0.169	0.168	0.164	0.316
	(0.016)	(0.039)	(0.016)	(0.019)	(0.016)	(0.043)
$\mathbf{1}\{\text{R\&D}>0\}_{it}$	0.221					
	(0.017)					
log (R&D/sales)$_{it}$		0.055				
		(0.007)				
$\mathbf{1}\{\text{R\&D+adv.}>0\}_{it}$			0.137			
			(0.011)			
log (R&D+adv.)/sales$_{it}$				0.036		
				(0.003)		
$\mathbf{1}\{\text{Intangibles}>0\}_{it}$					0.091	
					(0.017)	
log (Intangibles/sales)$_{it}$						0.008
						(0.008)
SIC6×year FE	✓	✓	✓	✓	✓	✓
Observations	118,579	19,649	118,579	75,879	118,579	16,051
R^2	0.043	0.120	0.041	0.058	0.038	0.100

Note: Specialization is the negative of vertical span, standardized within industry-year: positive coefficients indicate that the characteristic predicts more specialization. adv. = advertising; Intangibles = intangible assets. Each column is a separate regression. Standard errors in parentheses are clustered at the firm level.

A.4 Supplier Base: Effective Number of Partners and Link Duration

What does the supplier base of a specialized firm look like? We start by calculating the effective number of suppliers in the network, where we weight business partners by their transaction shares. Following Ekerdt and Wu (2024), we compute firm-level Effective Number of Customers (ENC) and Effective Number of Suppliers (ENS) as follows:

$$\text{ENC}_i = \frac{1}{\text{B2B Sales HHI}_i} = \left[\sum_{j \in \mathscr{I}_i^c} \left(\frac{x_{ji}}{\text{B2B Sales}_i} \right)^2 \right]^{-1}$$

$$\text{ENS}_i = \frac{1}{\text{B2B Purchases HHI}_i} = \left[\sum_{j \in \mathcal{J}_i^s} \left(\frac{x_{ij}}{\text{B2B Purchases}_i} \right)^2 \right]^{-1}$$

where $\mathcal{J}_i^c$ denotes firm i's set of customers and $\mathcal{J}_i^s$ the set of suppliers, and x_{ij} the transaction value from supplier j to firm i. The two measures are bounded above by the actual number of partners and decrease as transactions become more concentrated among specific partners.

For each supplier- or customer-side outcome y_{it} (raw count of partners, log effective number, or Herfindahl index), we estimate

$$y_{it} = \beta \cdot \text{Specialization}_{it} + \mathbf{X}_{it}'\delta + \mu_{k(i)t} + \varepsilon_{it}$$

where $\mathbf{X}_{it}$ contains log sales, log employment, and firm age, and $\mu_{k(i)t}$ is a SIC6-by-year fixed effect. Table A4 shows that, conditional on firm scale, specialized firms transact with more suppliers, have a higher effective number of suppliers, and exhibit lower sourcing concentration; the customer-side measures show no systematic relationship with specialization. Specialization is not synonymous with dependence on a narrow supplier base: a firm that concentrates on its core stage spreads its input purchases across a broader, less concentrated set of suppliers.

Table A4: Vertical Specialization and the Structure of the Supplier Base

	#Suppliers$_{it}$	#Customers$_{it}$	log ENS$_{it}$	log ENC$_{it}$	Sourcing HHI$_{it}$	Sales HHI$_{it}$
Specialization$_{it}$	3.936	-1.539	0.066	-0.005	-0.030	0.002
	(1.162)	(1.318)	(0.007)	(0.007)	(0.001)	(0.002)
SIC6×Year FE	✓	✓	✓	✓	✓	✓
Controls	log Sales, Age, log Employment					
Observations	118,579	118,579	118,562	116,535	118,562	116,535
R^2	0.581	0.087	0.273	0.103	0.206	0.099

Note: Specialization is the negative of vertical span, standardized within industry-year: positive coefficients indicate that more specialized firms exhibit more of the outcome. ENS (ENC) is the effective number of suppliers (customers), the inverse of the B2B purchase (sales) HHI. Sourcing (Sales) HHI is the Herfindahl index of B2B purchases (sales) across partners. Controls: log sales, firm age, log employment; SIC6×year fixed effects. Standard errors in parentheses are clustered at the firm level.

Sourcing concentration within input industry. The firm-level concentration measures above aggregate over a firm's entire supplier base across several upstream industries. Figure A1 examines this concentration within input industries. For each buyer i and supplier j,

we compute the share of j in i's total purchases from j's 6-digit industry $k(j)$,

$$s_{ij}^{k(j)} = \frac{x_{ij}}{\sum_{j' \in \mathcal{J}_i^s \cap k(j)} x_{ij'}}$$

in the 2017 cross-section of the cleaned network. The distribution is sharply bimodal: a large mass of marginal suppliers with shares near zero, and a spike of comparable size at exactly one, where the buyer sources an entire input industry from a single supplier. Single-sourcing within narrowly defined input markets is thus common in Taiwan, mirroring the pattern that Dhyne, Kikkawa, and Magerman (2022) document for Belgium. This is complementary to, not in tension with, the firm-level results: specialized firms spread purchases across many suppliers in total, yet within any single input industry the relationship is often exclusive. This bilateral granularity motivates the oligopolistic firm-to-firm extension to microfound markups that we discuss in the conclusion.

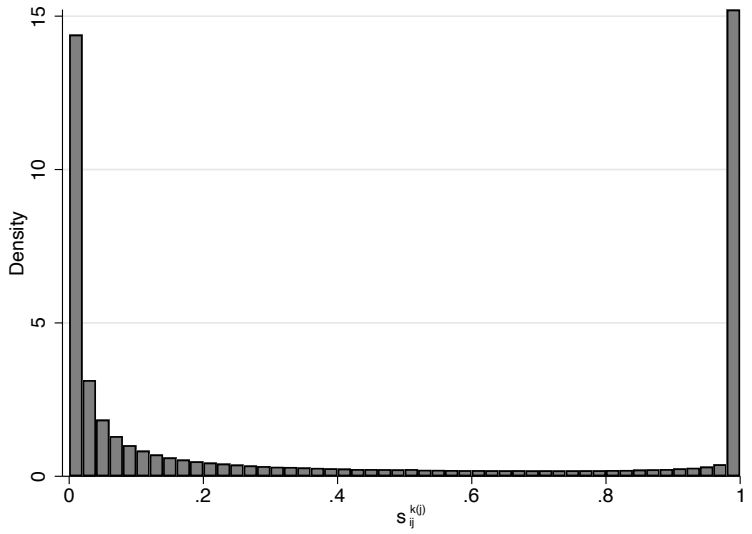


Figure A1: Distribution of Within-Industry Sourcing Shares

Note: Histogram of the pair-level sourcing share $s_{ij}^{k(j)}$, the share of supplier j in buyer i's total purchases from the supplier's 6-digit industry $k(j)$, across all buyer-supplier pairs in the 2017 cross-section of the cleaned network. A share of one indicates that the buyer sources that input industry from a single supplier.

Length of relationship with suppliers. We further examine the duration of linkages. For firms in the balanced panel, we calculate the transaction share-weighted duration of realized relationships (from start to termination), using the average transaction share of a buyer or seller. To alleviate concerns about censoring (buyer-seller relationships that start before 2013

and end after 2022), our baseline analysis focuses on links that start after 2013 and end before 2022.[24] We conduct one robustness check in which we keep the potentially censored links, and the results remain qualitatively unchanged. Columns 1–2 of Table A5 report the cross-sectional regression at the firm level,

$$\text{AvgDuration}_i = \beta \cdot \overline{\text{Specialization}}_i + \mathbf{X}_i'\delta + \mu_{k(i)} + \varepsilon_i$$

where $\overline{\text{Specialization}}_i$ is the firm's sample-period mean of $\text{Specialization}_{it}$ over years, $\mathbf{X}_i$ contains its mean log sales and log employment, and $\mu_{k(i)}$ is a SIC6 fixed effect. Columns 3–4 report the panel counterpart on year-to-year link survival rates,

$$\text{SurvivalRate}_{it} = \beta \cdot \text{Specialization}_{it} + \mathbf{X}_{it}'\delta + \mu_{k(i)t} + \varepsilon_{it}$$

with firm controls (log sales, log employment, and age) and SIC6-by-year fixed effects. Specialized firms have significantly longer relationships with their suppliers and do not exhibit a specific propensity to keep or dissolve links with customers; in the panel, the supplier links of specialized firms are also more likely to survive from one year to the next, while their customer links turn over slightly faster. Together with the supplier-base evidence above, the picture is one of *broad but sticky* sourcing: specialized firms spread purchases across more suppliers, likely across multiple industries, and keep those relationships longer.

A.5 Export Demand Shocks and Supplier Reallocation

We explore whether firms adjust their supplier networks in response to demand shocks by exploiting a shift-share export demand shock, in the spirit of several previous studies (Hummels et al. 2014; Huneeus 2018; Demir et al. 2024). To avoid potential confounding effects from the US-China trade war, we restrict the sample period to 2013–2017.

We define the quality-weighted demand shifter of HS6 product h from destination country c as

$$Z_{ch} = (\Delta \log \text{Imports}_{ch}) \times \log(\text{GDP per capita}_{c,2012})$$

where $\Delta \log \text{Imports}_{ch}$ is the import growth of product h in country c from all countries other than Taiwan, and $\log(\text{GDP per capita}_{c,2012})$ is the log GDP per capita for 2012 from the Penn

24. One type of uncensored link includes those that start after 2013 and end before 2022. The other type includes links in which either the buyer or seller enters after 2013, or exits before 2022. The tax data provide information about the establishment date and termination date of the business, allowing us to conduct this exercise more accurately.

Table A5: Vertical Specialization, Link Duration, and Link Survival

	Average duration (years)		Year-to-year survival rate	
	Suppliers	Customers	Suppliers	Customers
Specialization_i	0.130	-0.023		
	(0.029)	(0.033)		
$\text{Specialization}_{it}$			0.006	-0.008
			(0.001)	(0.001)
SIC6 FE	✓	✓		
SIC6×Year FE			✓	✓
Observations	4,209	4,209	104,990	104,990
R^2	0.089	0.045	0.131	0.070

Note: Columns 1–2: firm-level average link duration for the balanced panel, weighted by transaction shares; Specialization_i is the firm's sample-period mean of $\text{Specialization}_{it}$; Controls: log sales, log employment; heteroskedasticity-robust standard errors in parentheses. Columns 3–4: year-to-year transaction-share-weighted link survival rates in the full panel; Controls: log sales, firm age, log employment; standard errors in parentheses are clustered at the firm level.

World Table 10.01, serving as a proxy for the country's quality of demand as in Demir et al. (2024).[25]

The firm-level shock is

$$\text{ExportShock}_i = \sum_{c,h} x_{ich} \cdot Z_{ch}$$

where x_{ich} is the 2012 share of firm i's sales sold as product h to country c. The underlying assumption is that the import growth of product h in country c is positively correlated with country c's demand for product h from firm i, but is uncorrelated with other factors that increase firm i's demand. We then estimate

$$\begin{aligned}\Delta y_i \;=\; & \beta_1 \cdot \text{ExportShock}_i + \beta_2 \cdot \text{Specialization}_i \\ & + \beta_3 \cdot \left[\text{ExportShock}_i \times \text{Specialization}_i\right] + \mathbf{X}_i'\delta + \mu_{k(i)} + \varepsilon_i\end{aligned}$$

where Δy_i is the 2013–2017 long difference of a supplier- or customer-side quality moment ($\log \text{RDI}_{it}^{s,c}$, $\text{Avg}_{it}^{s,c}$, or $\text{Cov}_{it}^{s,c}$), $\mathbf{X}_i$ contains wage growth, sales growth, and the 2012 export-to-sales ratio (Borusyak, Hull, and Jaravel 2022), and $\mu_{k(i)}$ is a SIC6 fixed effect.

Table A6 presents the results. There is no direct effect of the export demand shock on the change in suppliers' R&D intensity; the interaction with specialization (third row) suggests

25. We use trade data from BACI to calculate the import growth, and we use *Asia, Not Elsewhere Specified* as a proxy for Taiwan's trade volume when excluding imports from Taiwan, as suggested by BACI itself.

that more specialized firms tilt their sourcing toward more R&D-intensive suppliers following the shock, driven by the covariance term. On the customer side, we observe no comparable adjustment.

Two caveats apply. Specialization enters here without the within-industry standardization used in the main text, so the interaction is largely driven by cross-industry variation in specialization. In an event-study specification, moreover, the same shocks show clear pre-trends. We therefore report this result as additional evidence rather than as a core result.

Table A6: Vertical Specialization, R&D Sorting, and Export Demand Shocks

	Supplier log RDI_{it}			Customer log RDI_{it}		
	ΔWeighted	ΔAvg	ΔCov	ΔWeighted	ΔAvg	ΔCov
$ExportShock_i$	0.011	0.009	0.002	-0.016	0.011	-0.027
	(0.017)	(0.011)	(0.016)	(0.016)	(0.012)	(0.012)
$Specialization_i$	−0.243	−0.198	−0.045	0.047	−0.031	0.079
	(0.126)	(0.079)	(0.106)	(0.119)	(0.085)	(0.092)
$Interaction_i$	0.088	−0.007	0.096	−0.008	0.001	−0.009
	(0.046)	(0.030)	(0.043)	(0.043)	(0.029)	(0.034)
SIC6 FE	✓	✓	✓	✓	✓	✓
Controls		Δlog wage, Δlog sales, $(X/\text{Sales})_{2012}$				
Observations	9,906	9,906	9,906	9,906	9,906	9,906
R^2	0.046	0.044	0.043	0.050	0.054	0.042

Note: The dependent variables are the differences between 2013 and 2017. $Specialization_i$ is the negative of the firm's average vertical span over 2013–2017, not standardized within industry. Controls: wage growth, sales growth, and the 2012 export-to-sales ratio. Heteroskedasticity-robust standard errors in parentheses.

A.6 Purchases from a Firm's Own Branches

Our specialization measure is meant to capture purchases from other firms. Built on the vertical span of production, it reflects specialization in production stages rather than the firm's legal boundaries.

The business registry records each entity's head office. In the VAT system, each business entity files under its own tax ID. A firm's branches with non-zero sales are required to file separately from its head office, and a purchase from a firm's own branch appears in the transaction data.

Using the head office information, we find that every firm in our analysis sample is its own head office. We then examine whether their suppliers report them as their head office, which informs us about within-firm sourcing. We classify a link as within-firm when the supplier lists

the buying firm as its head office. Table A7 shows that such links are rare: 655 to 1,664 pairs a year, or 0.01–0.02% of pairs. Their share of the sample firms' purchase value falls from 2.1% in 2013 to 0.5% in 2022 and has stayed below 1% since 2017. Only 1,660 firm-years, 1.4% of the sample, include any purchases from the firm's own branches.

Table A7: Purchases from a Firm's Own Branches

Year	Within-firm pairs	% of pairs	% of purchase value
2013	1,656	0.021	2.09
2014	1,664	0.021	1.24
2015	1,001	0.011	0.93
2016	952	0.011	1.01
2017	948	0.010	0.18
2018	823	0.009	0.69
2019	833	0.009	0.39
2020	916	0.010	0.64
2021	762	0.009	0.59
2022	655	0.008	0.53

Note: A within-firm pair is a buyer–supplier pair in which the supplier lists the buying firm, a firm in the analysis sample, as its head office. The last column is relative to the sample firms' total purchase value in the VAT records.

Table A8 regresses a firm's own-branch share of purchases on its specialization and R&D intensity, controlling for log sales, log employment, firm age, and industry-by-year fixed effects. Across all firm-years (Panel A) and among the 1,660 with positive own-branch purchases (Panel B), each coefficient has neither statistical nor economic significance. The links behind our specialization measure are therefore almost all between separate companies, and purchases from own branches do not vary systematically with specialization or R&D intensity. This is consistent with Atalay, Hortaçsu, and Syverson (2014), who find that upstream plants of vertically integrated U.S. firms ship little of their output to downstream plants of the same firm. Overall, such cases are too rare in Taiwan's industrial ecosystem to provide sufficient statistical power.

B Estimation Appendix

B.1 Solving for an Equilibrium

Holding fixed any set of parameters, an equilibrium of the model hinges on solving for the fixed point of the price vector $p^{\mathscr{U}}(q^u)$ and demand shifter $D^{\mathscr{U}}(q^u)$ for upstream firms, together with the sectoral price index p_s and aggregate income Y. The algorithm proceeds as follows:

1. Simulate $N = 10{,}000$ upstream and N downstream firms drawing ω_1^u and (ω_1^d, ω_1^h) from

Table A8: Own-Branch Share of Purchases, Specialization, and R&D Intensity

	Dependent variable: own-branch share of purchases		
	(1)	(2)	(3)
Panel A: all firm-years			
RDI_{it}	−0.0000		−0.0000
	(0.0000)		(0.0000)
$\text{Specialization}_{it}$		−0.0000	−0.0000
		(0.0003)	(0.0003)
log Sales_{it}	0.0002	0.0002	0.0002
	(0.0002)	(0.0002)	(0.0002)
Age_{it}	−0.0000	−0.0000	−0.0000
	(0.0000)	(0.0000)	(0.0000)
log Emp_{it}	0.0002	0.0002	0.0002
	(0.0001)	(0.0001)	(0.0001)
Observations		116,488	
R^2		0.007	
Panel B: firm-years with positive own-branch purchases			
RDI_{it}	−0.0095		0.0089
	(0.0097)		(0.0106)
$\text{Specialization}_{it}$		0.0065	0.0064
		(0.0135)	(0.0136)
log Sales_{it}	0.0076	0.0063	0.0072
	(0.0062)	(0.0059)	(0.0067)
Age_{it}	−0.0002	−0.0002	−0.0001
	(0.0005)	(0.0004)	(0.0004)
log Emp_{it}	−0.0129	−0.0113	−0.0123
	(0.0081)	(0.0077)	(0.0087)
Observations		1,660	
R^2	0.254	0.255	0.255
Industry×Year FE	✓	✓	✓

Note: The dependent variable is the share of a firm's domestic purchases bought from its own branches, suppliers that list the firm as their head office. RDI = R&D-to-sales; Specialization is the negative of vertical span, standardized within industry-year. Panel B keeps the firm-years with a positive own-branch share. Standard errors in parentheses.

the log-normal distributions, and an independent Hicks-neutral shifter $\omega_0 \sim \mathcal{N}(0, \sigma^2_{\omega_0})$ at each production node. Discretize the quality space into 46 evenly spaced points on $[0, 10]$, the first of which is the zero-quality option (no quality investment).

2. **Inner loop** (given Y): given guesses of the price vector $p^{\mathscr{U}}_{\text{old}}$ and demand vector $D^{\mathscr{U}}_{\text{old}}$ across the quality space, upstream firms find optimal q^u, yielding the updated price vector $p^{\mathscr{U}}_{\text{new}}$ according to Equation 2. Downstream firms find optimal (q^d, q^h) to maximize $\pi(d)$ according to Equation 13, computing the material price index according to Equation 8;

each firm draws $f_E \sim G_E$ and determines its export status by the inner maximization of Equation 19. With optimal downstream choices and realized export status, recalculate the demand shifter $D^{\mathscr{U}}_{\text{new}}$ according to Equation 15 and update the sectoral price index p_s according to Equation 11. Iterate until $|p^{\mathscr{U}}_{\text{new}} - p^{\mathscr{U}}_{\text{old}}| + |D^{\mathscr{U}}_{\text{new}} - D^{\mathscr{U}}_{\text{old}}|$ and the change in p_s fall below tolerance.

3. **Outer loop**: compute aggregate labor demand L_{demand} — production labor, quality investment, and export fixed costs — at the inner-loop solution. Update Y by bisection until $|L_{\text{demand}} - \bar{L}| < 10^{-3}\,\bar{L}$.

Because the wage is the numeraire, labor market clearing pins down aggregate income Y rather than w; for quintile moments that are scale-free, the normalization $\bar{L} = 1$ is without loss. It only helps pin down the scale-related parameters.

B.2 Moment Deviations

Let $g(\theta)$ stack the model-minus-data deviations across the 31 targeted moments. The estimator minimizes the GMM quadratic form

$$Q(\theta) = g(\theta)' \, \widehat{V}^{-1} g(\theta)$$

where $\widehat{V}$ is the variance–covariance matrix of the 31 data moments, estimated by a firm-clustered bootstrap of the analysis sample: each replication resamples firms with replacement and recomputes every moment, quintile and aggregate alike, and $\widehat{V}$ is the covariance of the moments across replications. Weighting by $\widehat{V}^{-1}$ is the efficient choice for this moment vector: precisely estimated moments receive more weight, correlated moments are not double-counted, and the objective does not depend on the units of any moment.

B.3 Standard Errors

Standard errors for the MSM estimates start from the GMM sandwich formula

$$\widehat{\text{Var}}(\hat{\theta}) = \left(\hat{G}' W \hat{G}\right)^{-1} \hat{G}' W \, \widehat{V} \, W \hat{G} \left(\hat{G}' W \hat{G}\right)^{-1}$$

which, because the weighting matrix is $W = \widehat{V}^{-1}$, collapses to $(\hat{G}' \widehat{V}^{-1} \hat{G})^{-1}$. Here $\hat{G}$ is the 31×13 Jacobian of the model moments with respect to the parameters, computed by central differences at the estimates with steps scaled to each parameter, and $\widehat{V}$ is the bootstrap variance–covariance matrix of the previous subsection, with a delta-method transform for the

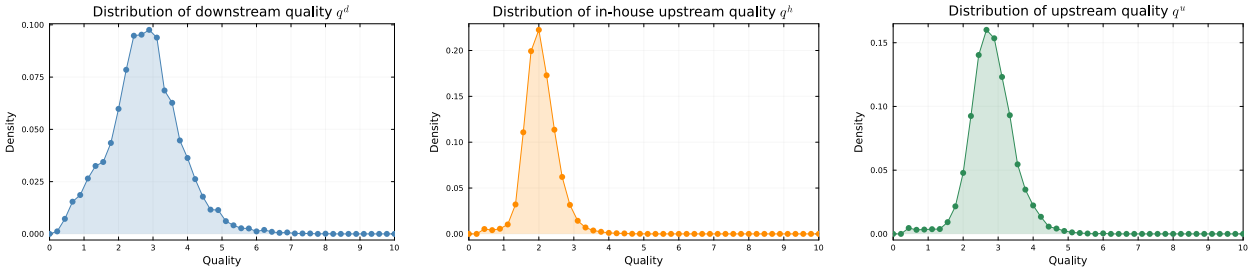
supplier-RDI profile normalized by its first-quintile value. The inversion is performed in elasticity units (columns of $\hat{G}$ scaled by the parameter values) for numerical conditioning.

C Additional Quantitative Results

C.1 Quality Distributions in the Estimated Baseline

This subsection describes the quality choices of the estimated model. Figure C1 plots the equilibrium distributions of the three quality objects: downstream output quality q^d, the in-house upstream quality q^h that downstream firms produce for themselves, and the quality q^u chosen by external upstream suppliers.

Compared with suppliers' quality (panel (c)), in-house upstream quality (panel (b)) is tightly concentrated: its distribution peaks near $q^h \approx 2$ and has essentially no mass above 4, because a downstream firm's in-house quality is limited by its own upstream capability ω_1^h. External suppliers, by contrast, choose quality that is both higher on average and more dispersed, with a right tail extending beyond $q \approx 4$. Panel (a) shows the distribution of output quality, whose long right tail is driven by the firms that source externally.



(a) Output quality q^d (b) In-house quality q^h (c) Supplier quality q^u

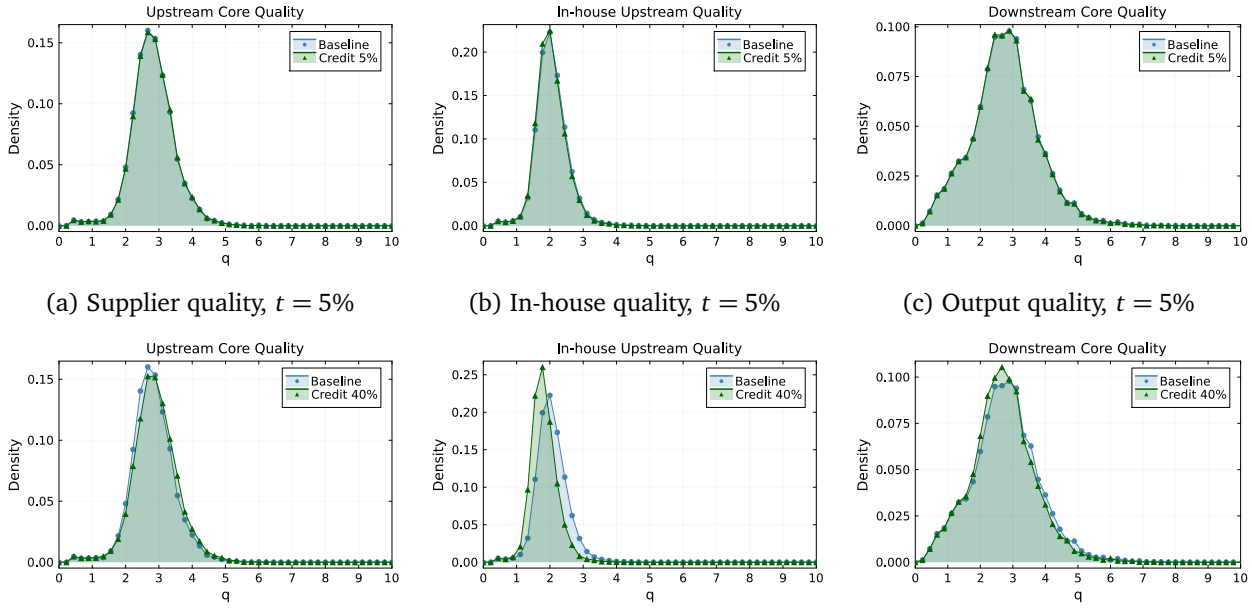
Figure C1: Quality Distributions in the Estimated Baseline

Note: Equilibrium distributions of chosen quality in the simulated baseline economy at the estimated parameters. Panel (a): output quality of downstream firms. Panel (b): quality of the in-house upstream variety chosen by downstream firms. Panel (c): quality chosen by external upstream suppliers. The quality space is discretized on $[0, 10]$.

C.2 Quality Distributions under the Credit

Figure C2 overlays the baseline and counterfactual distributions of quality under the two credits of Section 6.3. With a 5% credit, the distributions barely move, and the in-house share falls from 8.4% to 7.7%. With the full offset, the in-house share falls to 3.5%. Mean in-house quality

falls from 2.06 to 1.82, as firms make fewer of their inputs; mean supplier quality rises from 2.82 to 2.89, and its 90th percentile from 3.56 to 3.78. Mean output quality falls slightly, from 2.78 to 2.70, while its 90th percentile stays at 4.00.



(a) Supplier quality, $t = 5\%$ (b) In-house quality, $t = 5\%$ (c) Output quality, $t = 5\%$

(d) Supplier quality, $t^* = 40\%$ (e) In-house quality, $t^* = 40\%$ (f) Output quality, $t^* = 40\%$

Figure C2: Quality Distributions under a Credit on Input Purchases

Note: Distributions of chosen quality in the estimated baseline and under a credit that reimburses a fraction t of every external input bill, with real exports held at their baseline level (Table 8, columns 3 and 4).

D The Import Margin

This appendix extends the model of Section 4.2 with imported inputs and reports its estimation. The extension fits the model to the import moments: specialized firms are more likely to import, and, among importers, they spend a smaller share of their purchases on imports (Table D2). An estimated import margin is what a trade-policy counterfactual, such as a tariff on imported inputs or import substitution for the upstream sector, would require; we leave such exercises to future work. The extended model matches import participation and part of the import-share gradient, and the core estimates of Table 6 move little.

Foreign Supply. The international market supplies one foreign input variety of quality q_F, representing the average quality of inputs from the rest of the world. It is supplied at the world price parameterized as $b_3 q_F$, or $e\, b_3 q_F$ in home units. Importing requires a sunk cost in units of labor, $f_I \sim \mathcal{LN}(\mu_I, \sigma_I^2)$, drawn together with f_E after quality choices and independently across firms.

Sourcing with Imports. An importing firm combines its domestic purchases and the foreign variety into a purchases bundle Q_B with elasticity σ_F, and combines that bundle with its in-house variety at the elasticity σ of Equation 6:

$$m = \left[Q_B^{\frac{\sigma-1}{\sigma}} + [\eta \phi(q^d, q^h)]^{1/\sigma} x^h(d)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

$$Q_B = \left[B^{\frac{\sigma_F-1}{\sigma_F}} + I \left[\phi(q^d, q_F)^{\frac{1}{\sigma-1}} x^*(d) \right]^{\frac{\sigma_F-1}{\sigma_F}} \right]^{\frac{\sigma_F}{\sigma_F-1}}, \qquad B = \left[\int_{u \in \mathscr{U}} \phi(q^d, q^u)^{1/\sigma} x^u(d)^{\frac{\sigma-1}{\sigma}} \, du \right]^{\frac{\sigma}{\sigma-1}}, \tag{33}$$

where $I \in \{0,1\}$ is the import status and $x^*(d)$ the quantity imported. If $\sigma_F = \sigma$, the foreign variety enters exactly as a domestic supplier of quality q_F would. The same complementarity function ϕ governs every input. The foreign variety competes with domestic purchases at the elasticity of substitution σ_F rather than σ.

Let $P_{\text{ext}}^{1-\sigma} = \int_{\mathscr{U}} \phi(q^d, q^u)\, p^{\mathscr{U}}(u)^{1-\sigma}\, du$ denote the external term of Equation 8, and let

$$\tilde{p}_F(q^d) = \frac{e\, b_3 q_F}{\phi(q^d, q_F)^{1/(\sigma-1)}}$$

be the quality-adjusted price of the foreign variety to a firm of output quality q^d. The updated price indices are

$$P_B^{1-\sigma_F} = P_{\text{ext}}^{1-\sigma_F} + I\, \tilde{p}_F^{\,1-\sigma_F}, \qquad P_m^{1-\sigma} = \eta\, \phi(q^d, q^h)(p^h)^{1-\sigma} + P_B^{1-\sigma} \tag{34}$$

with $P_B = P_{\text{ext}}$ for a non-importer, in which case Equation 34 reduces to Equation 8. The import share of purchases is $s_F = \tilde{p}_F^{\,1-\sigma_F}/P_B^{1-\sigma_F}$, and the import share of all materials is $s_F\, P_B^{1-\sigma}/P_m^{1-\sigma}$.

The nest is what lets the model reproduce the import-share gradient, where specialized firms import a smaller share of what they buy. The intuition runs through complementarity. The estimated foreign variety sits below the domestic quality ladder, so it suits low-quality buyers better than high-quality ones: as a buyer's output quality rises, the foreign variety becomes a worse match for it, its quality-adjusted price rises relative to that of domestic suppliers, and the buyer shifts spending toward domestic upstream firms. If the foreign variety entered the same CES aggregate as domestic suppliers, this shift would be too small to match the data. The complementarity function bounds how far the foreign variety's weight can fall across buyers. The nest changes how a given gap in quality-adjusted prices translates into spending. Domestic and foreign purchases substitute at the elasticity σ_F rather than σ, so when the two are closer substitutes, the same price gap moves the expenditure share $(\sigma_F - 1)/(\sigma - 1)$ times as much. Import participation, by contrast, rises with firm size through the ability to cover the fixed cost.

The Firm's Problem with Two Statuses. The firm chooses its qualities before observing (f_E, f_I) and its two statuses after. Its problem extends Equation 19 to four status combinations:

$$\pi(d) = \max_{q^d, q^h} \mathbb{E}_{f_E, f_I}\left[\max_{(E,I)\in\{0,1\}^2} \sigma^{-1}\left(p^{\mathscr{D}}(I)\right)^{1-\sigma} D^{\mathscr{D}}(q^d, E) \;-\; wE\,f_E \;-\; wI\,f_I\right] - wf(q^d) - wf(q^h) \tag{35}$$

where $p^{\mathscr{D}}(I)$ is the price of Equation 9 evaluated at the material price index P_m of Equation 34 for the given import status. Importing lowers the firm's marginal cost; the gain is larger for firms with larger revenue, so exporters and large firms import more often, and larger for firms whose output quality is close to q_F.

Equilibrium. Three conditions of Section 4.2 change. In the upstream demand Equation 15, the share of a downstream firm's material spending that reaches domestic suppliers is $(1 - s_F)\,P_B^{1-\sigma}/P_m^{1-\sigma}$ rather than $P_{\text{ext}}^{1-\sigma}/P_m^{1-\sigma}$, with each domestic variety's share within it, $\phi(q^d, q^u)(p^{\mathscr{U}})^{1-\sigma}/P_{\text{ext}}^{1-\sigma}$, unchanged. Aggregate import spending is

$$M = \int_{\mathscr{D}} I(d)\,\frac{\alpha_m}{\mu}\,s_F(d)\,\frac{P_B(d)^{1-\sigma}}{P_m(d)^{1-\sigma}}\,r^{\mathscr{D}}(d)\,dd \tag{36}$$

since material spending is the share α_m of a firm's cost $r^{\mathscr{D}}/\mu$. The trade surplus $X - M$ replaces X in the goods-market clearing condition Equation 21, and the labor-market clearing condition Equation 22 gains the import fixed costs $\int_{\mathscr{D}} I(d) f_I(d)\,dd$. As in the main text, $e = 1$ in the

estimated equilibrium, and no stance on the current account is needed for estimation.

Parameters and Identification. The extension adds five parameters: the import fixed-cost distribution (μ_I, σ_I), the foreign price level b_3, the foreign quality q_F, and the elasticity σ_F. They are estimated jointly with the thirteen parameters of the main text on the 31 moments of Section 5 plus 13 import moments computed on the same cross-section: the share of importing firms and the mean import share of purchases by specialization quintile (demeaned by 6-digit industry, as the other quintile moments), and three aggregate levels, the importer share, the mean import share of purchases, and the mean ratio of imports to COGS. The mapping is direct. The importer share pins μ_I; the aggregate import intensity pins b_3 given q_F; the participation gradient across quintiles is informative for σ_I; and the import-intensity gradient is informative for σ_F, which scales how strongly the quality-adjusted foreign price moves the share. The level shifters b_3 and q_F are not separately identified when q_F is well below the domestic ladder, since both then act on the level of $\tilde{p}_F$; we report both but interpret only their product.

Estimates and Fit. Table D1 reports the estimates. The foreign variety is located below the domestic quality range ($q_F = 1.5$, against supplier qualities of roughly 2 to 5), priced at 0.75 in units of the home wage, and the nest elasticity is high ($\sigma_F = 8.4$), at the upper end of estimates of the elasticity between domestic and imported inputs. The fixed-cost dispersion is large ($\sigma_I = 8.8$): import participation is driven mostly by the cost draw rather than by the gain from importing, which is how the model reconciles a high importer share among small specialized firms with a low import share among them. The thirteen core parameters move little relative to Table 6; the largest changes are in the capability correlation (ϱ_ω from -0.41 to -0.38), the quality-cost curvature (γ from 7.8 to 6.9), and the in-house weight (η from 0.09 to 0.13), and the complementarity parameter ν is essentially unchanged.

Table D2 reports the import moments. Participation is matched in sign and size: importing falls by 17.5 percentage points from the most specialized to the most integrated quintile in the model, against 12.7 in the data. The import share of purchases is matched at the specialized end, where the model reproduces the low share of quintile 1 and the rise through quintile 3, but not at the integrated end: the model's share turns down in quintiles 4 and 5 where the data continue to rise. The shortfall is structural rather than numerical. In the model, integrated firms are the low-quality firms, and their in-house channel, outside the nest, does not affect their import share of purchases; what lowers their share is that their revenue and thus their gain from importing are small. In the data, integrated firms import a larger share of what they buy, which the single-variety, single-ladder specification has no channel to produce.

Table D1: Parameter Estimates with the Import Margin

Description	Parameter	Main text	With imports
Import block			
Mean of $\log f_I$	μ_I	—	−13.547
S.D. of $\log f_I$	σ_I	—	8.844
Foreign price level	b_3	—	0.500
Foreign variety quality	q_F	—	1.492
Purchases-bundle elasticity	σ_F	—	8.440
Core parameters			
Mean of $\log \omega_1$	μ_ω	0.679	0.640
S.D. of $\log \omega_1$	σ_ω	0.493	0.462
Quality curvature	$\bar{\omega}$	−0.292	−0.358
S.D. of Hicks-neutral shifter	σ_{ω_0}	0.336	0.347
Capability correlation	ϱ_ω	−0.411	−0.385
Quality-cost scale (10^{-9})	F	2.41	3.06
Quality-cost curvature	γ	7.780	6.885
Quality complementarity	ν	0.429	0.431
In-house weight	η	0.092	0.130
Foreign demand shifter (10^{-6})	b_1	1.84	2.32
Foreign demand curvature	b_2	3.067	3.048
Mean of $\log f_E$	μ_E	−11.737	−12.003
S.D. of $\log f_E$	σ_E	2.218	1.785

Note: Method of simulated moments estimates; 44 moments, 18 parameters. The "Main text" column reproduces Table 6. Standard errors for the extension have not been computed. The foreign price in units of the home wage is $b_3 q_F = 0.75$; b_3 and q_F are not separately identified (see text).

Table D2: Model Fit: Import Moments by Specialization Quintile

Moment		Q1 (specialized)	Q2	Q3	Q4	Q5 (integrated)
Share of importing firms	Data	0.031	0.056	0.023	−0.015	−0.096
	Model	0.089	0.004	0.014	−0.019	−0.087
Import share of purchases	Data	−0.033	−0.026	−0.011	0.014	0.056
	Model	−0.024	0.002	0.019	0.011	−0.009
Aggregate levels						
Importer share	Data 0.504	Model 0.549				
Import share of purchases	Data 0.192	Model 0.157				
Imports over COGS	Data 0.090	Model 0.094				

Note: Quintile moments are demeaned by 6-digit industry, as in Table 7; quintile 1 is the most specialized. The import share of purchases is imports over imports plus domestic firm-to-firm purchases, as in Figure 1.