

# Multiproduct Firms and Production Networks

Renata Abbott<sup>1</sup> and Marco Rojas<sup>2</sup>

<sup>1</sup>UC San Diego<sup>†</sup>

<sup>2</sup>Central Bank of Chile<sup>\*</sup>

October 2026

[Click here for the latest version.](#)

## Abstract

Do product scope and customer acquisition interact in shaping supplier size? Using the universe of firm-to-firm transactions in Chile, we show that they are interconnected. Within suppliers, a 10% increase in buyers is associated with a 5.6% increase in product scope. We develop a model of multiproduct suppliers in a production network in which the two margins are jointly determined through amortization of product-adoption costs across buyers and matching, whereby a broader buyer base creates more opportunities for each product. The model separates these channels and assesses their importance. Disciplining the model with the allocation of products across buyer relationships implies that matching accounts for most co-movement. Spillovers primarily run from buyers to products: buyer-acquisition frictions shape product scope, while product-adoption costs barely affect buyer networks. Finally, the buyer margin alone reproduces about 89% of firm-sales dispersion, compared with 53% for the product margin alone, with product scope amplifying differences generated through buyer networks.

**JEL Codes:** D21, D22, F12, F14, L11, L14, O31

**Keywords:** Multiproduct firms; Production networks; Firm-to-firm trade; Buyer–supplier relationships; Firm size dispersion; Customer acquisition; Product scope; Firm dynamics

---

<sup>†</sup>Email: [rabbottnavarrete@ucsd.edu](mailto:rabbottnavarrete@ucsd.edu). Website: [renataabbott.github.io](https://renataabbott.github.io).

We are grateful to Marc Muendler, Juan Herreño, Fabian Eckert, Munseob Lee, Gabriel Cañedo, Harrison Mitchell, and John Juneau for their helpful comments. We also thank participants at the CESifo Junior Workshop on International Economics, the Southern California IMF Conference, and internal seminars at the Central Bank of Chile and UC San Diego.

<sup>\*</sup>The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Central Bank of Chile or of anyone else associated with the Central Bank of Chile.

# 1 Introduction

How do suppliers' product expansion and buyer acquisition decisions interact within production networks, and how do they shape supplier growth and the firm-size distribution? Previous work demonstrates that these two margins are quantitatively important, but they are typically studied separately. Product introduction is an important determinant of supplier performance (Arkolakis et al., 2021), while buyer acquisition is an important determinant of supplier growth and the firm-size distribution (Bernard et al., 2022). When suppliers sell many products to many buyers, product expansion and buyer acquisition are separate decisions, but expanding along one might change the incentives to expand along the other.

We study these decisions using the universe of invoice-level firm-to-firm transactions in Chile (SII Factura Electrónica), matched to administrative firm characteristics.<sup>1</sup> We focus on manufacturing suppliers and define products at the 5-digit CPC level. We refer to the number of buyers served by a supplier as its buyer base. Observing which products each supplier sells to each buyer allows us to distinguish buyer acquisition from product expansion in determining supplier sales. We present three novel facts that motivate a structural model to quantify their interaction and assess its importance for the firm-size distribution.

First, the largest supplier-growth episodes disproportionately combine new buyers with new products. Decomposing sales growth between 2017 and 2023 according to whether buyers and products are new or already present shows that simultaneous expansion along both margins becomes especially important among the fastest-growing suppliers. Based on this fact we can conclude that these two margins frequently interact and shape firm sales jointly.

Second, the relative importance of these two margins depends on supplier size. Smaller suppliers rely relatively more on product expansion and buyer acquisition together, while larger suppliers increasingly expand by selling existing products to both new and existing buyers. Across the supplier size distribution, the contribution of buyers seems to be the largest determinant of sales growth.

Third, the buyer base and product scope co-move strongly within suppliers. In our preferred specification, a 10 percent increase in the number of buyers is associated with approximately a 5.6 percent increase in product scope. The relationship remains positive when product scope is measured only among incumbent buyers, with an elasticity of 0.38, approximately 69 percent of the baseline elasticity. The co-movement therefore extends beyond the additional products purchased by new buyers. This fact allows us to quantify the interaction between these margins and provides us with a key relationship for our quantitative model to match.

Guided by these facts, we build a static model in which suppliers jointly choose product scope and the buyers they serve. The two decisions are jointly determined through two different model

---

<sup>1</sup>Throughout the paper, whenever we refer to firms, we mean suppliers.

mechanisms: amortization and matching. Developing a product requires an adoption cost paid once by the supplier, while its returns accrue across buyer relationships. Access to a larger buyer base can therefore make it more profitable to develop additional products. The matching channel arises because a product's profitability varies across buyers, so a larger buyer network gives each product more opportunities to find a profitable match. A supplier with more buyers can therefore sell a larger share of the products it has already developed without expanding its developed portfolio. Conversely, a broader product portfolio increases the expected surplus from a potential buyer relationship, making buyer acquisition more valuable.

One potential concern about the interaction between the two margins is that cannibalization could limit the gains from product expansion and buyer acquisition, with new sales coming at the expense of existing sales. We present additional empirical evidence at both the firm and link level following product introductions and buyer acquisitions, and find little support for such effects in our setting. When a supplier introduces a new product within an existing buyer relationship, sales of its existing products to that buyer do not systematically decline over the subsequent two years. Similarly, when a supplier acquires a large new buyer, sales to existing buyers do not systematically fall. These patterns suggest that expansion along either margin increases the scale of supplier activity rather than merely redistributing a fixed amount of revenue across products or buyers.

Our data allow us to discipline the model not only with firm-level moments, but also with how products are allocated across individual buyer–supplier relationships. We show that this is crucial because two models can generate the same correlation between buyers and products at the firm level while implying very different responses to changes in network frictions. The calibrated model reproduces the relationship with an elasticity of 0.52, against 0.56 in the data. Across specifications, amortization contributes at most about 12 percent to the co-movement. This implies that matching determines the vast majority of the observed relationship between buyer base and product scope.

We then present two counterfactuals to quantify how buyer acquisition and product expansion contribute to the supplier size distribution, accounting for their co-movement and the spillovers between them. The first one shows that cutting buyer acquisition costs in half increases the buyer base by about 45 percent and product scope by about 14 percent. However, cutting product adoption costs in half raises product scope by about 7 percent but buyer base by only about 2 percent. Spillovers run mainly from buyers to products. This distinction is important when evaluating policies aimed at digitalization, trade integration, business connectivity, etc. An observed increase in the number of products sold does not necessarily mean that innovation increased; it may instead reflect improved matching inside the production network. Because the model is partial equilibrium, these responses do not include adjustments in buyers' price indices.

The second counterfactual quantifies which margin matters more for the firm-size distribution.

We find that the buyer base is the dominant one. Restricting suppliers to a single product shows that the buyer margin alone reproduces about 89 percent of the dispersion in supplier size generated by the full model, measured by the standard deviation of log sales. In contrast, replacing each supplier’s buyer base with a single representative buyer shows that the product margin alone reproduces 53 percent of this dispersion. Products still matter, but mainly by amplifying differences that originate in suppliers’ buyer bases.

The paper contributes to the production-network literature by modeling product scope as a distinct choice that interacts with buyer acquisition. This literature has developed rich models of buyer–supplier matching and network formation (Bernard et al., 2018, 2022; Eaton et al., 2026; Arkolakis et al., 2023a; Huneus, 2018a; Lim, 2018). However, these models typically treat suppliers as single-product or subsume product expansion into match formation. The single-product assumption is restrictive in our setting: 86% of suppliers in our data sell multiple products. Our model shows that the central role of the buyer base in explaining supplier-size heterogeneity, emphasized by Bernard et al. (2022), persists when suppliers are allowed to be multiproduct. Product scope contributes mainly by amplifying differences arising from suppliers’ buyer networks.

The paper contributes to the multiproduct literature in two ways. First, we study product expansion in supplier–buyer relationships within production networks. A large literature emphasizes the role of product introduction and product switching in firm performance and resource reallocation. Much of this work models final-consumer demand (Bernard et al., 2011; Mayer et al., 2014; Argente et al., 2018; Copeland and Shapiro, 2016). In this setting, Hottman et al. (2016) find that product scope makes a substantial contribution to firm-size heterogeneity, although demand appeal remains the dominant source. We complement this work by examining suppliers’ sales to other firms, where product expansion interacts with the acquisition of individual buyers.

Second, we allow buyer-specific product appeal to shape product selection within supplier–buyer relationships. Models in this literature generally abstract from heterogeneity across individual supplier–buyer relationships. In related work on international trade, Arkolakis et al. (2021) study how multiproduct exporters choose their product scope across destinations. Their model combines a common production-cost ladder with destination-specific appeal shocks. These shocks are realized after firms choose their product scope and therefore affect quantities sold within an already selected product portfolio.

Our empirical evidence shows that abstracting from heterogeneity at the buyer–product level misses important features of suppliers’ sales in this setting. Most suppliers sell multiple products, yet different buyers purchase different subsets of the same supplier’s portfolio. Moreover, products that rank highly in a supplier’s overall sales need not rank highly within each relationship. A common supplier-level product ranking is therefore insufficient to explain the allocation of products across buyers. To account for these patterns, we develop a buyer-specific product-matching structure that preserves a supplier-level product ranking while allowing product profitability to

vary across buyers. In our model, buyer-specific appeal is observed before the supplier decides which products to sell to each buyer, so it affects both the number and identity of products sold within each relationship. A supplier may develop many products, but each buyer purchases only the subset that is sufficiently profitable within that relationship. Product scope therefore depends on which buyers a supplier reaches, as well as on the products it develops.

Finally, our findings speak to the design of industrial policy, particularly in developing economies. Recent work highlights the importance of economic structure and production linkages for understanding policy effects (Juhász et al., 2024; Liu, 2019), while evidence of different responses among small and large firms underscores the relevance of firm heterogeneity (Criscuolo et al., 2019). We show that suppliers of different sizes also rely differently on product expansion and buyer acquisition: smaller suppliers rely relatively more on simultaneous expansion along both margins, whereas larger suppliers increasingly grow by selling existing products to new and incumbent buyers. Moreover, the dominant role of matching in our model implies that improved access to buyers can expand the range of products sold without requiring suppliers to develop additional products. This mechanism complements evidence that business-network interventions and relationships with multinational buyers can improve firm performance (Cai and Szeidl, 2018; Alfaro-Ureña et al., 2022).

Our results therefore suggest that support for product development and adoption should be considered jointly with measures that reduce buyer-search and relationship-formation costs, such as business-matching initiatives and programs connecting domestic suppliers to potential buyers. For manufacturing suppliers selling to other firms, a focus on innovation alone may overlook an important source of growth: finding profitable buyers for products they can already supply.

The remainder of the paper is organized as follows. Section 2 describes the data sources and sample construction and reports descriptive statistics. Section 3 presents the three key empirical facts. Section 4 develops and characterizes the model, while Section 5 describes its calibration. Section 6 assesses model fit and presents an untargeted validation exercise based on the composition of supplier growth. Section 7 presents the mechanism decomposition and counterfactual analyses. Section 8 concludes.

## 2 Data Description

### 2.1 Administrative data sources

The main dataset combines three administrative sources. Our core data consist of domestic firm-to-firm transactions among Chilean firms, drawn from tax records collected by the Chilean tax authority, the Servicio de Impuestos Internos (SII). These records cover the universe of formal private-sector firms. The SII began collecting these data in 2014, and electronic invoicing became

mandatory in 2017. Our sample contains invoices dated from January 2015 through December 2023 and is accessed through the Central Bank of Chile.<sup>2</sup>

Each invoice has two parts. The first contains transaction-level information: anonymized tax identifiers for the supplier and buyer, the value of the transaction, the municipality associated with each party (corresponding to the municipality of the firm’s registered headquarters), and payment terms (cash versus trade credit, including maturity). The second part contains line-item details on the goods and services traded, including a free-text product description, price and quantity (including discounts or surcharges when applicable), and a product identifier. Because the product identifier is assigned at the level of unique text descriptions, the same good is often split into multiple IDs due to minor differences in wording, and the IDs can be extremely granular.<sup>3</sup> We therefore harmonize products using the Central Product Classification (CPC)<sup>4</sup>.

We merge the transaction records with two additional administrative datasets. The first contains firm balance-sheet and tax information from annual form F22 and monthly form F29, including total sales, intermediate-input expenditures, the wage bill, and fixed assets. The second is the employer–employee dataset from form DJ1887, which reports annual earnings paid by each firm to each worker and months of employment. Reported earnings include wages, salaries, bonuses, tips, and other taxable labor-income components. All datasets are anonymized.

## 2.2 Units of observation and network definitions

Throughout the paper, a firm is defined by a unique anonymized tax identifier. Unless stated otherwise, the term firm refers to a supplier. Buyers and suppliers can belong to any sector of the Chilean economy, although the main stylized facts focus on manufacturing suppliers. Suppliers include both upstream (primary) and downstream (secondary) firms in the network, i.e., any firm that sells a good or service to another firm. Transactions are strictly business-to-business because purchases by final consumers are not observed.

Let  $s$  denote a supplier,  $b$  a buyer,  $g$  a 5-digit CPC product, and  $t$  a year. Let  $x_{sbg t}$  denote

---

<sup>2</sup>This paper was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from public and private entities through collaboration agreements. To protect the privacy of workers and firms, the CBC requires that the development, extraction, and publication of results not permit the direct or indirect identification of individuals or firms. CBC officials processed the disaggregated data. All analysis was implemented by the authors and does not involve or compromise the SII. Information in the SII databases is of a tax nature and originates in taxpayer self-declarations; its veracity is therefore not the responsibility of the SII.

<sup>3</sup>For example, “BOY SNEAKER” and “GIRL SNEAKER” are classified as different products.

<sup>4</sup>The CPC is an internationally harmonized classification system that organizes goods and services based on standardized concepts, definitions, and coding rules.

sales from supplier  $s$  to buyer  $b$  of product  $g$  in year  $t$ . We define

$$\begin{aligned}\mathcal{B}_{st} &= \{b : \exists g \text{ such that } x_{sbgt} > 0\}, \\ \mathcal{G}_{st} &= \{g : \exists b \text{ such that } x_{sbgt} > 0\}, \\ \mathcal{R}_{st} &= \{(b, g) : x_{sbgt} > 0\}.\end{aligned}$$

The corresponding counts are

$$B_{st} = |\mathcal{B}_{st}|, \quad G_{st} = |\mathcal{G}_{st}|, \quad R_{st} = |\mathcal{R}_{st}|.$$

Here,  $B_{st}$  is the number of buyers (buyer base),  $G_{st}$  is the number of products sold (product scope), and  $R_{st}$  is the number of active buyer–product pairs.

Supplier sales observed in the transaction network are

$$S_{st}^{\text{B2B}} = \sum_b \sum_g x_{sbgt}.$$

When total supplier sales are available from tax filings, we denote them by  $S_{st}^{\text{total}}$ .

The corresponding extensive-margin counts are  $B_{st}$ , the supplier’s buyer base;  $G_{st}$ , its product scope; and  $R_{st}$ , the number of active buyer–product pairs. The average number of buyers per product is therefore

$$\frac{R_{st}}{G_{st}},$$

while the average number of products per buyer is

$$\frac{R_{st}}{B_{st}}.$$

We also use the ratio

$$\frac{B_{st}}{G_{st}},$$

which measures the size of the buyer base relative to product scope.

Sales within a supplier–buyer link are

$$S_{sbt}^{\text{link}} = \sum_g x_{sbgt},$$

while sales associated with a supplier–product relationship are

$$S_{sgt}^{\text{prod}} = \sum_b x_{sbgt}.$$

## 2.3 Sample construction

We apply several cleaning and sample restrictions. First, we retain firms that report more than one paid worker. Second, we restrict attention to firms that file monthly tax form F29 in all sample years used in the relevant exercise and report positive sales. Following [Burstein et al. \(2024\)](#), we remove transactions that exceed USD 125 million, transactions with non-positive prices or quantities, and invoices in which the supplier and buyer are the same firm. Finally, we use product-level minimum and maximum price filters from the Chilean National Statistics Institute (INE) to remove line items with implausible unit prices. The exact estimation sample varies across exercises because different results require different levels of aggregation and support.

## 2.4 Descriptive statistics

Table 1 reports descriptive statistics for manufacturing suppliers in 2017 and 2023, separately for multi-product and single-product suppliers. Multi-product production is pervasive in the Chilean manufacturing sector. Multi-product suppliers account for 86.2% of suppliers in 2017 and 83.8% in 2023, and they operate at substantially greater scale than single-product suppliers. In 2017, the average multi-product supplier has total sales of \$2.35 million and serves 68.9 buyers, compared with \$0.97 million in sales and 7.6 buyers for the average single-product supplier. These differences remain large in 2023. This suggests that the interaction between product scope and the buyer base is relevant for most manufacturing suppliers and is especially important among larger suppliers.

Multi-product suppliers also exhibit substantial heterogeneity in both dimensions of scope. In 2017, they sell an average of 16.1 products and serve 68.9 buyers, while the median supplier sells only nine products and serves nine buyers. Similar gaps between means and medians remain in 2023 and are especially large for the number of buyers. Sales are also highly right-skewed. This means that a relatively small group of suppliers operate with very broad product portfolios and large buyer bases. Capturing this heterogeneity across suppliers is an important goal of the quantitative model.

## 2.5 Product scope, buyer base, and supplier size

Figure 1 documents two familiar relationships between supplier size and the buyer base and product scope. Panel (a) shows that suppliers with higher sales serve more buyers, while Panel (b) shows that larger multiproduct suppliers sell more products. These patterns are consistent with previous evidence linking firm size to the buyer base ([Bernard et al., 2022](#); [Arkolakis et al., 2023b](#)) and product scope ([Hottman et al., 2016](#); [Bernard et al., 2010](#); [Argente et al., 2018](#)). Multiproduct suppliers also tend to serve more buyers than single-product suppliers at comparable sales levels. Both margins are therefore closely associated with supplier size. However, these supplier-level

Table 1: Descriptive Statistics for Chilean Manufacturing Suppliers

|                                           | 2017   |        |           | 2023   |        |           |
|-------------------------------------------|--------|--------|-----------|--------|--------|-----------|
|                                           | Mean   | Median | Std. Dev. | Mean   | Median | Std. Dev. |
| <i>Panel A: Multiproduct suppliers</i>    |        |        |           |        |        |           |
| Total sales (million USD, supplier level) | 2.35   | 0.08   | 43.95     | 4.22   | 0.10   | 84.29     |
| B2B intermediate sales (million USD)      | 0.96   | 0.02   | 29.53     | 1.67   | 0.02   | 53.283    |
| Number of buyers                          | 68.904 | 9.000  | 784.554   | 59.914 | 8.000  | 667.048   |
| Number of products                        | 16.061 | 9.000  | 20.955    | 15.764 | 8.000  | 21.583    |
| Number of buyers per product              | 3.749  | 1.000  | 22.768    | 3.220  | 1.000  | 20.357    |
| Product specialization ratio              | 0.541  | 0.512  | 0.243     | 0.553  | 0.528  | 0.245     |
| Buyer concentration ratio                 | 0.514  | 0.475  | 0.287     | 0.544  | 0.514  | 0.286     |
| Number of suppliers                       | 37,179 | –      | –         | 28,938 | –      | –         |
| Share of all suppliers (%)                | 86.2   | –      | –         | 83.8   | –      | –         |
| <i>Panel B: Single-product suppliers</i>  |        |        |           |        |        |           |
| Total sales (million USD, supplier level) | 0.971  | 0.022  | 25.378    | 1.166  | 0.023  | 27.427    |
| B2B intermediate sales (million USD)      | 0.268  | 0.002  | 10.445    | 0.272  | 0.002  | 6.572     |
| Number of buyers                          | 7.638  | 1.000  | 32.583    | 4.767  | 1.000  | 22.324    |
| Number of products                        | 1.000  | 1.000  | 0.000     | 1.000  | 1.000  | 0.000     |
| Number of buyers per product              | 7.638  | 1.000  | 32.583    | 4.767  | 1.000  | 22.324    |
| Product specialization ratio              | 1.000  | 1.000  | 0.000     | 1.000  | 1.000  | 0.000     |
| Buyer concentration ratio                 | 0.801  | 1.000  | 0.295     | 0.845  | 1.000  | 0.259     |
| Number of suppliers                       | 5,935  | –      | –         | 5,585  | –      | –         |
| Share of all suppliers (%)                | 13.8   | –      | –         | 16.2   | –      | –         |


*Notes:* Sales are in million USD. Shares are computed within year among single-product and multiproduct suppliers. A supplier is classified as multi-product in a year if it sells at least two distinct 5-digit CPC products in that year. The product-specialization ratio is the share of supplier sales accounted for by its largest product. The buyer-concentration ratio is defined analogously using the supplier’s largest buyer. Buyers per product are calculated at the supplier level before averaging across suppliers.

counts leave open how products are distributed across buyers, which is central to understanding the relationship between the two margins.

Our data allow us to examine this allocation directly. Figure 2 relates supplier size to two additional measures: the average number of buyers purchasing each product and the average number of products purchased by each buyer. Panel (a) shows that larger suppliers sell each product to more buyers on average. Panel (b) shows that their buyers also purchase more products from them on average. Larger suppliers thus differ along both aspects of the buyer–product structure: their products reach a broader set of buyers, and their buyer relationships encompass broader product baskets. The increase in buyers per product is particularly pronounced, whereas products per buyer rises more gradually and levels off among the largest suppliers. Figure 10 in Appendix A further documents this asymmetry using the ratio of the buyer base to product scope.

These additional patterns motivate studying buyer acquisition and product expansion jointly. Each product can generate sales across multiple buyer relationships, while each relationship can generate sales from multiple products. Acquiring another buyer creates a potential outlet for the supplier’s products, and offering additional products can increase what the supplier sells within a buyer relationship. The returns to either decision can therefore depend on the other. Under-

Figure 1: Extensive margins and supplier size



*Notes:* The figure presents binned scatterplots with fitted regression lines relating supplier size to the number of buyers and products. Supplier size is measured by total sales at the tax-identifier level (million USD). Both axes are shown on a log scale. The sample excludes the top 0.1 percent of the sales distribution. Each dot represents a bin of suppliers based on total sales. Panel (a) distinguishes multiproduct and single-product suppliers, while Panel (b) reports multiproduct suppliers only, since single-product suppliers have exactly one product by definition.


standing supplier size requires accounting for both the breadth of suppliers' product portfolios and buyer bases and the allocation of products across individual relationships. This motivates a framework that keeps product expansion and buyer acquisition as distinct choices while allowing them to interact. The next section examines this interaction through the composition of supplier growth, how this composition varies across the firm-size distribution, and the within-supplier co-movement of buyers and products.

Separate accounting decompositions provide further evidence that both margins are relevant, with a larger contribution from the buyer base. In 2017, the number of buyers accounts for 42.9 percent of the variance in log supplier sales, while average sales per buyer account for the remaining 57.1 percent. In the corresponding product decomposition, the number of products accounts for 27.3 percent of the variance, while average sales per product account for 72.7 percent. Figure 11 in Appendix A.1 reports these results. The buyer and product shares are alternative decompositions of the same variation in supplier sales and should not be added across decompositions.

### 3 Empirical Facts

This section documents three model-facing facts about how buyer base and product scope interact within suppliers. First, the largest supplier-growth episodes disproportionately combine the addition of new buyers and new products. Second, the composition of growth varies systematically with suppliers' initial size: smaller suppliers rely relatively more on margins involving new

Figure 2: Intensive margins and supplier size



*Notes:* The figure presents binned scatterplots relating supplier size to the density of buyer–product relationships. Panel (a) reports the average number of buyers per product, defined as the number of active buyer–product pairs divided by the number of products. Panel (b) reports the average number of products per buyer, defined as the number of active buyer–product pairs divided by the number of buyers. Supplier size is measured by observed firm-to-firm sales (million USD). The horizontal axis is shown on a log scale. Each dot represents one of 50 equally populated sales bins, and the solid line is a quadratic fit. The top 1 percent of the sales distribution is excluded.

products, while larger suppliers increasingly expand by selling established products to new buyers and through existing buyer–product relationships. Third, buyer base and product scope co-move strongly within suppliers, providing direct evidence on the interaction between the two margins.

The first two facts characterize how suppliers move between an initial and a final allocation and therefore have a natural comparative-static counterpart in the model. The third fact disciplines the model’s central mechanism linking buyer base and product scope.

Appendix B provides complementary evidence that supplier–buyer links and supplier–product relationships are empirically distinct objects. Link sales peak earlier and decline faster, cohort profiles show a more systematic deterioration in sales with link age, and survival rates are consistently lower for supplier–buyer links than for supplier–product relationships. Because our model is static, these patterns are not quantitative targets of the model. Rather, they motivate treating buyer base and product scope as separate margins of adjustment.

### 3.1 Fact 1: Supplier growth combines buyer and product expansion

We begin by studying how suppliers move between an initial and a final allocation. Although this decomposition is measured over time, its model counterpart is a comparative-static exercise: the static model can compare a baseline equilibrium with a counterfactual equilibrium and decompose the resulting change across buyer and product margins.

Let  $t \in \{2017, 2023\}$ . Define the sets of buyers and products of supplier  $s$  in the baseline year:

$$\mathcal{B}_{s0} = \{b : \exists g \text{ such that } x_{sb,2017} > 0\}, \quad \mathcal{G}_{s0} = \{g : \exists b \text{ such that } x_{sb,2017} > 0\}.$$

We classify each unit of 2023 sales according to whether its buyer and product were present in 2017:

$$\begin{aligned} S_{s,2023}^{\text{OLOP}} &= \sum_{b \in \mathcal{B}_{s0}} \sum_{g \in \mathcal{G}_{s0}} x_{sb,2023}, & \text{old link, old product,} \\ S_{s,2023}^{\text{NLOP}} &= \sum_{b \notin \mathcal{B}_{s0}} \sum_{g \in \mathcal{G}_{s0}} x_{sb,2023}, & \text{new link, old product,} \\ S_{s,2023}^{\text{OLNP}} &= \sum_{b \in \mathcal{B}_{s0}} \sum_{g \notin \mathcal{G}_{s0}} x_{sb,2023}, & \text{old link, new product,} \\ S_{s,2023}^{\text{NLNP}} &= \sum_{b \notin \mathcal{B}_{s0}} \sum_{g \notin \mathcal{G}_{s0}} x_{sb,2023}, & \text{new link, new product.} \end{aligned}$$

Because all 2017 sales occur with buyers and products that are old relative to the 2017 baseline,

$$S_{s,2017} = S_{s,2017}^{\text{OLOP}}.$$

The four components of supplier sales growth are therefore:

$$\begin{aligned} \Delta S_s^{\text{OLOP}} &= S_{s,2023}^{\text{OLOP}} - S_{s,2017}, \\ \Delta S_s^{\text{NLOP}} &= S_{s,2023}^{\text{NLOP}}, \\ \Delta S_s^{\text{OLNP}} &= S_{s,2023}^{\text{OLNP}}, \\ \Delta S_s^{\text{NLNP}} &= S_{s,2023}^{\text{NLNP}}. \end{aligned}$$

They satisfy the accounting identity

$$\Delta S_s = S_{s,2023} - S_{s,2017} = \Delta S_s^{\text{OLOP}} + \Delta S_s^{\text{NLOP}} + \Delta S_s^{\text{OLNP}} + \Delta S_s^{\text{NLNP}}. \quad (1)$$

The OLOP component is the net change in sales within the supplier's baseline buyer-product space. It includes sales changes on continuing baseline buyer-product pairs, losses from baseline pairs that disappear, and sales on newly formed pairs combining a buyer and a product that were each present in 2017. Thus, "old buyer-old product" does not necessarily mean that the specific buyer-product pair traded in the baseline year. The other three components involve at least one buyer or product that is new relative to the baseline.

**Decomposition by realized growth.** We first group suppliers by their realized 2017–2023 sales growth and compute the average contribution of each component to  $\Delta S_s$ . Figure 3 asks

an ex post accounting question: among suppliers that subsequently shrink or grow by different amounts, which margins account for their realized change in sales?

The lower tail displays little offsetting expansion through new buyers or new products. Contraction is concentrated in OLOP, reflecting lower sales in the continuing portion of the supplier’s baseline buyer–product space and losses from exiting relationships. Because each component is normalized by the supplier’s total change in sales, the OLOP share appears positive for shrinking suppliers: both the OLOP contribution and total sales growth are negative, so their ratio is positive. By contrast, any positive contribution from new buyers or new products is divided by a negative total change and therefore appears as a negative share. In the upper tail, extensive margins account for a much larger share of growth. The NLNP component shows that the fastest-growing suppliers disproportionately combine new buyers and new products rather than expanding along only one margin.


The increasing importance of NLNP among the fastest-growing suppliers is also visible when we isolate the three margins involving a new buyer or a new product. Appendix C.1 and C.2 provide two additional decompositions that reinforce these patterns. The first one isolates the three margins involving a new buyer or a new product by excluding OLOP and renormalizing NLOP, OLNP, and NLNP relative to their net extensive-margin adjustment. The composition shifts strongly toward NLNP as realized supplier growth increases, with the joint new-buyer–new-product margin accounting for the largest share among the fastest-growing suppliers. Second, appendix C.2 separates OLOP into changes among baseline buyer–product relationships that remain active through 2023 and losses from relationships that disappear. Relationship destruction accounts for most of the contraction in the lower tail, whereas the largest expansions rely increasingly on the formation of new buyer–product combinations. Together, these additional decompositions reinforce the main result that large supplier expansions disproportionately involve simultaneous growth along the buyer and product margins.

### **3.2 Fact 2: As suppliers grow larger and older, expansion shifts from new products and buyers toward existing products and new buyers**

The previous figures classify suppliers according to their realized change in sales. The sorting variable is therefore the outcome being decomposed. We now ask a different question: do suppliers that begin from different initial conditions expand through different combinations of buyer acquisition, product introduction, and sales within continuing relationships? We examine two dimensions of initial heterogeneity: suppliers’ position in the 2017 sales distribution and their stage in the firm life cycle.

**Heterogeneity by initial supplier scale.** We first assign suppliers to groups according to their total sales in 2017 and then decompose their sales growth between 2017 and 2023. This

Figure 3: Decomposition of supplier sales growth, 2017–2023



*Notes:* The figure decomposes supplier sales growth between 2017 and 2023 across intervals of the realized-growth distribution. Each bar reports the average contribution of new link–new product (NLNP), new link–old product (NLOP), old link–new product (OLNP), and old link–old product (OLOP). Each category corresponds to the indicated interval of the realized-growth distribution.


exercise restricts attention to suppliers with positive growth. Accordingly, the figure describes the composition of growth among expanding suppliers; it does not establish which initial-size groups grow faster or are more likely to grow.

Figure 4 reveals a clear pattern across the initial-size distribution. Among suppliers with positive growth, initially smaller suppliers rely relatively more on margins involving new products, particularly the joint addition of new buyers and new products (NLNP). As initial supplier size increases, growth shifts toward existing products: larger suppliers rely relatively more on sales through existing buyer–product relationships (OLOP) and on forming new buyer links for products already in their portfolios (NLOP). Thus, product expansion is relatively more important for initially smaller suppliers, while larger suppliers expand relatively more by deepening existing relationships and extending established products to new buyers.

This increasing importance of the buyer margin among larger suppliers is consistent with its broader role in accounting for cross-sectional differences in firm size. In 2017, a decomposition of the variance of log supplier sales into the number of buyers and average sales per buyer attributes 42.9 percent to the buyer extensive margin and 57.1 percent to sales per buyer. The analogous product-side decomposition attributes only 27.3 percent to the number of products, with the

remaining 72.7 percent accounted for by sales per product. Thus, although both dimensions of scope are quantitatively important, differences in customer reach account for a larger share of supplier-size dispersion than differences in product scope. Appendix Figure 11 reports both decompositions.

Figure 4: Composition of positive sales growth by initial supplier size



*Notes:* Suppliers are grouped according to their position in the 2017 total-sales distribution. Within each initial-size group, the figure decomposes positive supplier sales growth between 2017 and 2023 into new link–new product (NLNP), new link–old product (NLOP), old link–new product (OLNP), and old link–old product (OLOP). Because the sample is restricted to suppliers with positive growth, the figure describes how expanding suppliers grow rather than differences in average growth rates or growth probabilities across initial-size groups. Each category corresponds to the interval between adjacent reported percentile cutoffs.

The distinction between Figures 3 and 4 is important. The first is an ex post accounting exercise: it asks which margins characterize different realized growth outcomes. The second is an ex ante heterogeneity exercise that asks whether the composition of subsequent expansion varies with a supplier characteristic measured before that expansion occurs. The latter is therefore especially useful for disciplining heterogeneous comparative-static responses in the model.

This heterogeneity has a direct model counterpart. Initial supplier sales can reflect base-line differences in productivity, buyer base, or product scope. The model can therefore apply a common change in the economic environment and ask whether initially small and large suppliers respond through different combinations of buyer acquisition, product adoption, and changes within continuing relationships. The cross-sectional variance decomposition provides a complementary benchmark for assessing the relative quantitative importance of the buyer and product margins for the firm-size distribution.

Appendix C.3 provides an additional view of the initial-size gradient by grouping suppliers into broad low-, middle-, and high-sales categories. It also reports the decomposition after excluding

OLOP, which focuses on growth coming from the extensive margins. The same pattern remains. The contribution of NLNP declines with initial supplier size, while NLOP becomes increasingly important. Even among transactions involving a new buyer or a new product, smaller suppliers are more likely to expand both margins together, whereas larger suppliers rely relatively more on selling products already in their portfolios to new customers.

Appendix C.4 examines the second dimension of initial heterogeneity, supplier age, while conditioning on sales in 2017. The life-cycle decomposition shows a similar pattern. Younger suppliers rely relatively more on NLNP, while NLOP becomes more important as firms age. More mature suppliers increasingly expand by selling products already in their portfolios to new buyers. Because age is not an explicit state variable in our static model, we treat this evidence as complementary to the initial-size results rather than as a separate quantitative target of the model.

### 3.3 Fact 3: buyer base and product scope co-move within suppliers

The previous facts establish that buyer relationships and product scope are distinct margins. We now ask whether these two margins are systematically related within suppliers. A priori, their relationship could reflect either a trade-off, whereby expanding along one margin comes at the expense of the other, or a complementarity, whereby a larger buyer base raises the return to maintaining a broader product portfolio. A broader product portfolio, in turn, raises the value of forming additional buyer relationships because it gives the supplier more products through which to generate surplus with a potential buyer, making the fixed cost of forming the relationship easier to cover.

This question is closely related to recent work on the interaction between demand-expansion investments, such as advertising, and innovation. For example, [Cavenaile and Roldan-Blanco \(2021\)](#) provide evidence of substitution between R&D and advertising among U.S. firms, whereas recent models often assume complementarity between customer expansion and innovation, as in [Einav et al. \(2021\)](#) and [Afrouzi et al. \(2023\)](#). These studies focus primarily on final-consumer settings. Our setting instead allows us to study product and customer scope directly within a production network.

We begin by relating a supplier’s product scope to its buyer base. Let  $\text{Buyers}_{st}$  denote the number of distinct buyers served by supplier  $s$  in year  $t$ , and let  $\text{Products}_{st}$  denote the number of distinct products (CPCs) it sells. We estimate:

$$\log(\text{Products}_{st}) = \beta_1 \log(\text{Buyers}_{st}) + \theta \log(\text{Sales}_{st}) + \gamma_t + \delta_s + \varepsilon_{st}, \quad (2)$$

where  $\gamma_t$  are year fixed effects and  $\delta_s$  are supplier fixed effects, respectively. The supplier fixed effects are particularly important for our purposes because they absorb persistent differences across firms that may jointly shape buyer base and product scope. In the model below, these

persistent firm characteristics are supplier productivity and the networking cost shifter. Thus, the preferred specification compares the same supplier over time, holding fixed its underlying type. Contemporaneous sales are included as an additional control for the supplier’s current scale. The coefficient  $\beta$  therefore captures whether, conditional on persistent supplier heterogeneity and current firm size, years in which a supplier serves more buyers are also years in which it operates a broader product portfolio. We interpret this relationship as a conditional association rather than a causal effect of buyer acquisition on product scope.

Table 2 reports the results. The relationship is strong even within suppliers. With only year and supplier fixed effects, the elasticity of product scope with respect to buyer base is 0.555, which is the empirical moment that maps most closely into the model because the model is static and therefore cannot reproduce within-supplier changes over time literally. Controlling additionally for supplier sales reduces the coefficient only modestly, to 0.486. Because sales are themselves an outcome of the two margins, we treat the specification without the sales control (column 5) as the benchmark target for the model.

Appendix D repeats the analysis with product scope measured only among incumbent buyers. The elasticity is 0.382, compared with 0.555 in the baseline specification; with the sales control, the corresponding coefficients are 0.288 and 0.486. Thus, the positive association persists within pre-existing buyer relationships and is not driven solely by products purchased by newly acquired buyers. However, this exercise does not distinguish newly developed products from products already developed by the supplier but newly sold to incumbent buyers. The ratio of the estimated elasticities is therefore not an identified share of co-movement attributable to product development or amortization. Section 7.1 compares the incumbent-buyer association with its model counterpart.

The positive relationship is not specific to the stock of products. Appendix E considers product additions, defined as products sold by the supplier in year  $t$  that were not sold in  $t - 1$ . With year and supplier fixed effects, a 10 percent increase in buyers is associated with approximately 3.92% additional product introductions. Thus, within suppliers, years with more buyers are also associated with more product introductions. The magnitude is economically meaningful.

Taken together, these results establish a robust positive relationship between the two dimensions of firm scope. Firms with broader customer bases also operate broader product portfolios, and years in which a supplier serves more buyers are associated with more product additions. These patterns do not by themselves identify a structural complementarity, since unobserved firm-level shocks may affect both margins. Instead, we treat their systematic co-movement as a model-facing fact. The evidence therefore motivates a framework in which buyer acquisition and product adoption are jointly determined rather than independent choices. In the model, the two margins are linked through two channels: the organization of fixed costs (amortization) and buyer-specific product matching. A broader buyer base raises the return to developing products and gives existing products more chances to find a buyer, while a broader product portfolio in-

creases the value of forming buyer relationships. We then ask whether these mechanisms can account for the empirical co-movement between buyer base and product scope, together with some of the other facts documented above.

Table 2: Buyer Breadth and Product Scope

| <i>Dependent variable: Log Number of Products</i> |                       |                        |                       |                        |                       |                       |
|---------------------------------------------------|-----------------------|------------------------|-----------------------|------------------------|-----------------------|-----------------------|
|                                                   | (1)                   | (2)                    | (3)                   | (4)                    | (5)                   | (6)                   |
| Log Number of Buyers                              | 0.498***<br>(0.00157) | 0.464***<br>(0.00183)  | 0.518***<br>(0.00147) | 0.458***<br>(0.00179)  | 0.555***<br>(0.00324) | 0.486***<br>(0.00352) |
| Log Sales                                         |                       | 0.0487***<br>(0.00135) |                       | 0.0877***<br>(0.00139) |                       | 0.127***<br>(0.00258) |
| Observations                                      | 138,784               | 138,534                | 138,784               | 138,534                | 138,219               | 137,966               |
| Firms                                             | 23,382                | 23,380                 | 23,382                | 23,380                 | 22,817                | 22,812                |
| $R^2$                                             | 0.434                 | 0.439                  | 0.565                 | 0.579                  | 0.886                 | 0.890                 |
| Year FE                                           | Yes                   | Yes                    | Yes                   | Yes                    | Yes                   | Yes                   |
| Sector FE                                         | No                    | No                     | Yes                   | Yes                    | No                    | No                    |
| Supplier FE                                       | No                    | No                     | No                    | No                     | Yes                   | Yes                   |

*Notes:* The dependent variable is the log number of distinct products (CPCs) sold by supplier  $s$  in year  $t$ . The main explanatory variable is the log number of distinct buyers served by the supplier in the same year. Columns (1)–(2) include year fixed effects; columns (3)–(4) additionally include supplier-sector fixed effects; and columns (5)–(6) include supplier fixed effects instead. Even-numbered columns additionally control for contemporaneous log supplier sales. Column (5) is the model’s calibration target; column (6) is reported for robustness. Heteroskedasticity-robust standard errors are reported in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

## 4 Model

The empirical analysis shows that firm size in production networks reflects two distinct extensive margins: product scope and buyer relationships. These margins are strongly interconnected within firms, and their interaction matters for the firm-size distribution because differences in scope can amplify underlying firm heterogeneity and generate substantial sales dispersion. Since both margins are endogenous and respond to common profitability forces, reduced-form correlations alone do not identify the mechanism linking them. We therefore develop a tractable static model in which suppliers jointly choose product scope and buyer relationships. The model serves three purposes. First, it isolates the mechanism generating complementarity between the two margins. Second, it quantifies how shocks to product-adoption or buyer-linkage costs propagate across both margins and reshape the firm-size distribution. Third, it assesses their contribution to firm-size dispersion while accounting for this complementarity.

Our model builds on the heterogeneous-firm logic of [Hopenhayn \(1992\)](#) and the rank-based multiproduct structure of [Arkolakis et al. \(2021\)](#), adapting them to a production-network en-

vironment in which destinations are individual buyers. The link-formation margin follows the literature modeling network formation through fixed costs per relationship (Lim, 2018; Huneus, 2018b; Bernard et al., 2022), rather than through search (Miyauchi, 2024; Arkolakis et al., 2023a). Suppliers are heterogeneous in productivity and choose how many products to develop and which buyers to serve. Product-adoption costs are paid once at the firm level and can therefore be amortized across buyer relationships, whereas link formation requires buyer-specific fixed costs. This distinction generates the model’s first source of complementarity, amortization. A larger buyer base increases the return to developing additional products, while a broader product portfolio makes forming new buyer relationships more valuable. Buyer-specific product matching creates a second source of complementarity, which we describe below.

The benchmark abstracts from displacement of existing sales as suppliers expand product or buyer scope. Appendix F documents substantial net expansion of supplier–buyer relationships and limited immediate displacement following product introduction, alongside declines in incumbent-product sales and retention at longer horizons. We therefore treat separability as a simplifying benchmark and consider within-link cannibalization separately in the nested-CES extension. A central modeling choice is how a supplier’s products map to individual buyers. Our baseline allows this mapping to be buyer-specific. Firms develop products at the supplier level, but their relevance for a particular buyer depends on an idiosyncratic product–buyer match. This choice is disciplined by the data. Firm-level and within-link product rankings are positively but imperfectly aligned, a pattern difficult to reconcile with a rigid common product ladder.<sup>5</sup> The matching structure also connects our framework to models of two-sided heterogeneity in firm-to-firm trade (Bernard et al., 2018; Eaton et al., 2026; Demir et al., 2024; Miyauchi, 2024), in which trading relationships and their intensity are endogenous.

## 4.1 Setting

Suppliers are indexed by  $s \in \mathcal{S}$  and buyers by  $b \in \mathcal{B}$ .<sup>6</sup> Supplier  $s$  has productivity  $\phi_s > 0$  and chooses a firm-level portfolio of  $G_s$  products. Supplier productivity follows a Pareto distribution,

$$\phi_s \sim \text{Pareto}(\theta_\phi, \phi_{\min}), \quad \phi_s \geq \phi_{\min} > 0. \quad (3)$$

Products are ordered from core to peripheral by rank  $g \in \{1, \dots, G_s\}$ , with product-level productivity

$$\phi_{s,g} = \frac{\phi_s}{h(g)}, \quad h(g) = g^\alpha, \quad \alpha \geq 0. \quad (4)$$

Products farther from the core are therefore less productive on average.

---

<sup>5</sup>A common product ladder, as in Arkolakis et al. (2021), imposes a shared ranking of products across buyers, so that buyers differ in how far down the same product ranking they purchase.

<sup>6</sup>We interpret suppliers as atomless from the perspective of each buyer’s input bundle. The finite supplier population used in the quantitative implementation approximates this underlying environment.

For each developed product  $g$  and potential buyer  $b$ , supplier  $s$  draws an idiosyncratic product–buyer match

$$\zeta_{sb} \sim \text{Pareto}(\theta_\zeta, \zeta_{\min}), \quad \theta_\zeta > 1, \quad \zeta_{\min} = \frac{\theta_\zeta - 1}{\theta_\zeta}, \quad (5)$$

so that  $\mathbb{E}[\zeta_{sb}] = 1$ . Product rank captures systematic differences in profitability across the supplier’s portfolio, while  $\zeta_{sb}$  allows product relevance to vary across buyers. Conditional on an active relationship, the set of products sold to buyer  $b$  is therefore buyer-specific,

$$\mathcal{G}_{sb} \subseteq \{1, \dots, G_s\}, \quad G_{sb} = |\mathcal{G}_{sb}| \leq G_s, \quad (6)$$

and need not form a common prefix of the supplier’s product ranking.

**Fixed costs and the amortization mechanism.** The supplier faces three fixed costs that operate at different levels of its problem and govern distinct margins of adjustment. First, developing product  $g$  requires the firm-level cost

$$f^P(g) = f^P(1)g^\delta, \quad \delta \geq 0, \quad (7)$$

so that a supplier developing  $G_s$  products pays total product-development costs

$$F^P(G_s) \equiv \sum_{g=1}^{G_s} f^P(g) = f^P(1) \sum_{g=1}^{G_s} g^\delta. \quad (8)$$

Crucially, this cost is paid once at the firm level. Once product  $g$  has been developed, it can potentially generate profits across all buyers served by the supplier.

Second, forming relationship  $(s, b)$  requires paying the link-level cost  $f_{sb}^L$ . Third, conditional on having developed product  $g$  and formed link  $(s, b)$ , selling that product to buyer  $b$  requires the product–buyer activation cost

$$f_{sb}^A(g) = f_{sb,0}^A g^\mu. \quad (9)$$

If  $\mathcal{G}_{sb} \subseteq \{1, \dots, G_s\}$  denotes the set of developed products that supplier  $s$  activates for buyer  $b$ , total activation costs on the link are

$$F_{sb}^A(\mathcal{G}_{sb}) \equiv \sum_{g \in \mathcal{G}_{sb}} f_{sb}^A(g) = f_{sb,0}^A \sum_{g \in \mathcal{G}_{sb}} g^\mu. \quad (10)$$

Thus, if  $\mathcal{B}_s$  denotes the set of buyers served by supplier  $s$ , its total fixed costs are

$$F_s^{\text{fixed}} = F^P(G_s) + \sum_{b \in \mathcal{B}_s} [f_{sb}^L + F_{sb}^A(\mathcal{G}_{sb})]. \quad (11)$$

This distinction across levels is central to the model’s complementarity between buyer base and product scope. Product-development costs are incurred once by the supplier and can therefore be amortized across buyer relationships. As the supplier serves more buyers, a newly developed product has more potential relationships over which to generate surplus, increasing the return to product adoption. Link and activation costs, by contrast, are not amortized across buyers: acquiring an additional buyer requires paying a new  $f_{sb}^L$ , and selling a particular product on that relationship requires paying the corresponding  $f_{sb}^A(g)$ . The firm-level location of  $F^P(G_s)$  is therefore essential for the model’s amortization force.

The complementarity also operates in the opposite direction. A supplier with a broader developed product portfolio can generate surplus from more potential product–buyer matches, raising the value of forming an additional buyer relationship. Buyer base and product scope are consequently jointly determined. Additional buyers increase the return to product development, while additional products increase the return to link formation. This extends the standard product-ladder structure of [Arkolakis et al. \(2021\)](#) to a production-network environment in which individual buyer relationships constitute the relevant destinations.

In the quantitative model, heterogeneity in link and activation costs is parameterized as

$$f_{sb}^L = \bar{f}_b^L \eta_{sb} c_s, \quad f_{sb,0}^A = \bar{f}_b^A \eta_{sb} c_s, \quad (12)$$

where  $\eta_{sb} > 0$  captures pair-specific cost heterogeneity and  $c_s > 0$  is the supplier’s networking cost, which captures persistent supplier heterogeneity in the cost of forming and operating buyer relationships. We normalize  $\mathbb{E}[\eta_{sb}] = \mathbb{E}[c_s] = 1$ . The supplier component  $c_s$  is common across all relationships of supplier  $s$ , while  $\eta_{sb}$  varies across supplier–buyer pairs. The pair-level cost shock  $\eta_{sb}$  and the product-level match shock  $\zeta_{sbg}$  capture distinct sources of heterogeneity.  $\eta_{sb}$  makes some supplier–buyer relationships systematically easier to establish and operate, while  $\zeta_{sbg}$  determines the suitability of an individual developed product for a particular buyer.

**Demand and operating profits.** We work in partial equilibrium. Buyer  $b$  has CES demand across available supplier–product varieties with elasticity  $\sigma > 1$ . Let  $E_b$  denote buyer expenditure and  $P_b$  its CES price index. Under monopolistic competition, product-level revenue is

$$r_{sbg} = D_b \phi_s^{\sigma-1} g^{-\alpha(\sigma-1)} \zeta_{sbg}, \quad (13)$$

where

$$D_b \equiv E_b \left( \frac{P_b}{[\sigma/(\sigma-1)]w} \right)^{\sigma-1} \quad (14)$$

summarizes buyer profitability. Variable operating profits are a constant share of revenue:

$$\pi_{sb} = \frac{1}{\sigma} r_{sb}. \quad (15)$$

The same match shock governs both selection into product–buyer trade and realized sales conditional on activation. Thus, unlike the deterministic-ladder benchmark, the matching specification does not require a separate ex-post product appeal shock to generate residual within-link sales dispersion. Appendix ?? derives these expressions from the buyer’s CES problem.

## 4.2 Supplier’s two-stage problem

The supplier’s problem has a recursive two-stage structure, which we solve by backward induction in expected value. In Stage 1, supplier  $s$  chooses its firm-level product scope  $G_s$  and pays the corresponding product-development costs. This determines the set of products that the supplier is capable of selling, but not the subset ultimately sold to each buyer.

In Stage 2, conditional on  $G_s$ , the supplier evaluates each potential buyer relationship. Link formation is an ex-ante decision. The supplier compares the expected surplus generated by its  $G_s$  developed products on relationship  $(s, b)$  with the buyer-specific link cost  $f_{sb}^L$ . If the relationship is formed, the product–buyer match shocks  $\zeta_{sb}$  are realized. Each developed product is then activated for buyer  $b$  whenever its realized operating profit covers its product–buyer activation cost  $f_{sb}^A(g)$ . Thus, firm-level scope  $G_s$  places an upper bound on link-level product scope, while the realized matches determine which subset of the developed portfolio is actually traded with each buyer.

Anticipating these Stage 2 decisions, the supplier chooses  $G_s$  in Stage 1 to maximize expected profits across all potential buyer relationships net of firm-level development costs. Figure 5 summarizes the sequence of decisions and uncertainty.

## 4.3 Optimality conditions

Fix firm-level scope  $G_s$ . Define the supplier–buyer profitability shifter

$$A_{sb} \equiv \frac{D_b}{\sigma} \phi_s^{\sigma-1}, \quad (16)$$


and let

$$\kappa \equiv \mu + \alpha(\sigma - 1). \quad (17)$$

Developed product  $g$  is profitable on link  $(s, b)$  whenever its realized match quality is sufficiently high:

$$\zeta_{sb} \geq z_{sb}(g) \equiv \frac{f_{sb,0}^A}{A_{sb}} g^\kappa. \quad (18)$$

Figure 5: Timing of the supplier's two-stage problem



*Notes:* In Stage 1, the supplier chooses firm-level product scope  $G_s$  and pays  $F^P(G_s)$ , anticipating the buyer-level problem. In Stage 2, buyer profitability  $D_b$  and pair-specific cost heterogeneity  $\eta_{sb}$  determine the expected value of each potential relationship. The supplier forms the link when expected product surplus covers  $f_{sb}^L$ . Conditional on an active link, product-buyer match shocks  $\zeta_{sbg}$  determine which developed products are activated and sold to that buyer.

The threshold  $z_{sb}(g)$  is the minimum match quality required to activate product  $g$  for buyer  $b$ . It is lower for more productive suppliers and more profitable buyers, making products easier to activate on those relationships. When  $\kappa > 0$ , the threshold also rises with product rank, so more peripheral products require stronger product–buyer matches to be traded.

Given the Pareto distribution of match quality, the corresponding probability that developed product  $g$  is activated for buyer  $b$  is

$$p_{sb}(g) \equiv \Pr(\zeta_{sb} \geq z_{sb}(g)) = \begin{cases} 1, & \text{if } z_{sb}(g) \leq \zeta_{\min}, \\ \left(\frac{\zeta_{\min}}{z_{sb}(g)}\right)^{\theta_\zeta}, & \text{if } z_{sb}(g) > \zeta_{\min}. \end{cases} \quad (19)$$

Because  $\zeta_{\min}$  is the lowest possible realization of match quality, a product whose activation threshold lies below this bound is activated with probability one. Otherwise, activation occurs only when the realized match exceeds the threshold, with probability given by the Pareto tail.

The parameter  $\theta_\zeta$  governs the dispersion of product–buyer match quality and therefore how the supplier’s firm-level portfolio maps into buyer-specific product sets. Conditional on firm-level scope  $G_s$ , expected product scope on link  $(s, b)$  is

$$\mathbb{E}[G_{sb} \mid G_s] = \sum_{g=1}^{G_s} p_{sb}(g). \quad (20)$$

Thus,  $G_s$  determines the set of products that the supplier can potentially sell, while product–buyer matches determine which subset of that portfolio is actually activated for each buyer.

Let expected net surplus from developed product  $g$  for buyer  $b$  be

$$\Psi_{sb}(g) \equiv \mathbb{E}_\zeta \left[ (A_{sb} g^{-\alpha(\sigma-1)} \zeta - f_{sb,0}^A g^\mu)^+ \right]. \quad (21)$$

Appendix I reports the corresponding closed-form Pareto expression. Expected gross surplus from potential relationship  $(s, b)$  is

$$C_{sb}(G_s) \equiv \sum_{g=1}^{G_s} \Psi_{sb}(g). \quad (22)$$

The supplier forms the relationship whenever this surplus covers the link cost:

$$I_{sb}^*(G_s) = \mathbf{1} \{C_{sb}(G_s) \geq f_{sb}^L\}. \quad (23)$$

The corresponding active buyer set is

$$\mathcal{B}_s(G_s) = \{b \in \mathcal{B} : C_{sb}(G_s) \geq f_{sb}^L\}, \quad N_s(G_s) = |\mathcal{B}_s(G_s)|. \quad (24)$$

Anticipating these buyer-level decisions, the supplier chooses firm-level scope to solve

$$G_s^* \in \arg \max_{G_s \in \{0, \dots, G_{\max}\}} \left\{ \sum_{b \in \mathcal{B}} [C_{sb}(G_s) - f_{sb}^L]^+ - F^P(G_s) \right\}, \quad (25)$$

where

$$F^P(G_s) = \sum_{g=1}^{G_s} f^P(g). \quad (26)$$

Conditional on an active relationship and realized matches, the set of products sold to buyer  $b$  is

$$\mathcal{G}_{sb} = \{g \leq G_s : \zeta_{sbg} \geq z_{sb}(g)\}, \quad G_{sb} = |\mathcal{G}_{sb}|. \quad (27)$$

Product sets can therefore differ across buyers and need not be nested.

#### 4.4 Equilibrium definition

Because the benchmark is partial equilibrium, buyer demand shifters  $\{D_b\}_{b \in \mathcal{B}}$  and the distribution of potential buyers are taken as given. Therefore, given model primitives and buyer demand shifters  $\{D_b\}_{b \in \mathcal{B}}$ , an equilibrium is a collection

$$\{G_s^*, \mathcal{B}_s^*, \mathcal{G}_{sb}^*, r_{sbg}^*\}_{s,b}$$

such that: (i)  $G_s^*$  solves (25); (ii) buyer relationships satisfy the link-formation rule in (23); (iii) products within active links satisfy the activation rule in (27); and (iv) product-level sales satisfy (13). For each candidate  $G_s$ , the link-formation rule is already embedded in the objective in (25). The equilibrium can therefore be computed directly by evaluating expected profits over the finite set of possible product scopes and selecting the maximizing value. Appendix J provides additional details.

#### 4.5 Characterization

The model's central implication follows directly from the different levels at which the fixed costs are paid. Product scope and buyer base reinforce one another in both directions.

**More products imply weakly more buyers.** Because  $\Psi_{sb}(g) \geq 0$ ,

$$C_{sb}(G_s + 1) = C_{sb}(G_s) + \Psi_{sb}(G_s + 1) \geq C_{sb}(G_s). \quad (28)$$

Hence,

$$\mathcal{B}_s(G_s) \subseteq \mathcal{B}_s(G_s + 1). \quad (29)$$

A broader product portfolio therefore weakly enlarges the set of profitable buyer relationships.

**A broader buyer base raises the return to product adoption.** Define the marginal value of adding product  $g$  as

$$\Delta\Pi_s(g) \equiv \Pi_s(g) - \Pi_s(g-1).$$

It can be decomposed as

$$\Delta\Pi_s(g) = \underbrace{\sum_{b \in \mathcal{B}_s(g-1)} \Psi_{sb}(g)}_{\text{surplus on existing links}} + \underbrace{\sum_{b \in \mathcal{B}_s(g) \setminus \mathcal{B}_s(g-1)} [C_{sb}(g) - f_{sb}^L]}_{\text{newly profitable buyer links}} - f^P(g). \quad (30)$$

The first term makes the amortization mechanism transparent. Since  $f^P(g)$  is paid once at the firm level, every additional active buyer adds another nonnegative expected-surplus term to the return from developing product  $g$ . The second term reinforces this effect because the new product may also make additional buyer relationships profitable.

A simple analytical approximation makes the first channel especially clear. In the interior Pareto region,

$$\Psi_{sb}(g) = \psi_{sb} g^{\mu - \kappa \theta \zeta}, \quad (31)$$

where

$$\psi_{sb} = \frac{\zeta_{\min}^{\theta \zeta}}{\theta \zeta - 1} A_{sb}^{\theta \zeta} (f_{sb,0}^A)^{1 - \theta \zeta}. \quad (32)$$

Holding the active buyer set locally fixed and abstracting from the discrete gain from newly activated links, optimal product scope is approximately

$$G_s^* \approx \left[ \frac{\Lambda_s^\psi}{f^P(1)} \right]^{\frac{1}{\delta + \kappa \theta \zeta - \mu}}, \quad \Lambda_s^\psi \equiv \sum_{b \in \mathcal{B}_s} \psi_{sb} = N_s \bar{\psi}_s, \quad \bar{\psi}_s \equiv \frac{1}{N_s} \sum_{b \in \mathcal{B}_s} \psi_{sb}. \quad (33)$$

This expression separates buyer base from average buyer value. Holding  $\bar{\psi}_s$  fixed, a larger buyer base  $N_s$  increases the return to the marginal product and therefore raises optimal product scope.

Equivalently, defining

$$\Lambda_s^M \equiv \sum_{b \in \mathcal{B}_s} D_b^{\theta \zeta} (f_{sb,0}^A)^{1 - \theta \zeta}, \quad (34)$$

we obtain

$$G_s^* \propto \left[ \frac{\phi_s^{\theta \zeta (\sigma - 1)} \Lambda_s^M}{f^P(1)} \right]^{\frac{1}{\delta + \kappa \theta \zeta - \mu}}. \quad (35)$$

Thus, conditional on supplier productivity, a more valuable buyer base raises optimal product scope. Together with the monotonicity of buyer-link formation in  $G_s$  established above, this

delivers the model’s two-way interaction between buyer base and product scope. The closed-form expression is used for economic interpretation; the quantitative model is solved using the exact discrete problem in equation (25).

**A second channel: matching.** Amortization operates on developed scope  $G_s$ . What the data record, however, is observed scope: the number of developed products sold to at least one buyer. Given  $G_s$  and the active buyer set, expected observed scope is

$$\mathbb{E}[\widehat{G}_s \mid G_s, \mathcal{B}_s] = \sum_{g=1}^{G_s} \left[ 1 - \prod_{b \in \mathcal{B}_s} (1 - p_{sb}(g)) \right], \quad (36)$$

which is increasing in the buyer set even when  $G_s$  is held fixed. A supplier with more buyers gives each developed product more independent chances to clear its activation threshold, so a larger share of its portfolio is observed in trade. This matching channel works entirely through selection. It does not require firms to develop additional products or pay any additional firm-level cost. Its strength is governed by how tightly selection binds, which depends on the activation cost  $\bar{f}_0^A$  and the dispersion of match quality  $\theta_\zeta$ . When activation almost never binds,  $p_{sb}(g) \approx 1$  and observed scope equals developed scope, so all co-movement comes from amortization. When selection binds, matching can dominate. Section 7.1 shows how to separate the two channels quantitatively and that the within-link moments, which discipline  $\bar{f}_0^A$  and  $\theta_\zeta$ , determine their relative importance.

The same decisions map directly into firm size. Conditional on realized matches, link sales and supplier sales are

$$t_{sb} = \sum_{g \in \mathcal{G}_{sb}} r_{sb}g, \quad T_s = \sum_{b \in \mathcal{B}_s} t_{sb}. \quad (37)$$

Heterogeneity in productivity, buyer composition, and the two endogenous scope margins therefore jointly shapes the firm-size distribution.

## 4.6 Estimation strategy

We discipline the quantitative model using firm-level distributional moments from the 2017 cross section of Chilean manufacturing suppliers, supplemented by within-link allocation statistics and the scope–buyer association estimated from the transaction panel. We calibrate the main specification by simulated method of moments (SMM). For each candidate parameter vector, we solve the supplier’s problem exactly, simulate the resulting product and buyer networks, and compare simulated moments with their empirical counterparts.

The targeted moments discipline the product and buyer margins separately, as well as their interaction. They include the firm-level distributions of product scope and buyer base, buyers per product, product specialization, buyer concentration, and the tail of the sales distribution. They also capture the scope–buyer relationship conditional on supplier productivity and the allocation

of products across buyer relationships, including the alignment between firm-level and within-link product importance. Rank-decay and rank-alignment moments are particularly informative about the systematic product ranking and the dispersion of product–buyer match quality, while the levels and dispersion of product scope and buyer base discipline the corresponding fixed-cost parameters and cost heterogeneity.

We target 17 moments with 10 free parameters. Eleven are firm-level moments: the median and the 75th and 90th percentiles of product scope and buyer base, the median number of buyers per product, buyer concentration, product specialization, the scope–buyer elasticity, and the sales tail index. The other six describe how products are allocated across a supplier’s buyers: within-link alignment, products per link, the share of single-product links, rank decay, the within-supplier dispersion of link sales, and portfolio coverage. The adoption-cost convexity  $\delta$  is estimated, and the activation-cost level is profiled rather than estimated jointly (Section 5.2). Parameters that are fixed externally or normalized are held constant throughout. We minimize the equally weighted sum of squared relative deviations between simulated and empirical moments. Moments outside the objective, including the mean and standard deviation of product scope, mean buyer base, and the single-product share, serve as validation checks. Appendix L reports a calibration that fixes  $\delta$  and compares the matching model with a common product ladder.

This strategy exploits information not only on firms’ total product and buyer scope, but also on how products are allocated across individual buyer relationships. The latter moments discipline the matching structure separately from the overall size of the two extensive margins. After estimation, we evaluate model fit and use the calibrated model to quantify the contribution of product scope, buyer base, and their complementarity to the firm-size distribution.

## 5 Model Calibration

### 5.1 Calibration strategy and empirical targets

We calibrate the model to the cross-sectional distribution of manufacturing suppliers in 2017. We use 2017 as the baseline year because electronic invoicing was mandatory by that year and because it provides a common cross section in which buyer relationships, product scope, and supplier sales can all be measured consistently. Using 2017 also allows us to avoid contamination from large aggregate shocks that affected the Chilean economy in subsequent years, including the social unrest that began in late 2019 and the COVID-19 pandemic. These events generated unusually large disruptions to firm sales, buyer–supplier relationships, and production decisions that are outside the mechanisms captured by our static model.

The calibration is designed to discipline the joint distribution of the two extensive margins at the center of the model, buyer base and product scope, as well as the allocation of products and

sales across those margins.

The quantitative analysis focuses on the matching model developed in Section 4. The calibration targets both firm-level moments and within-link moments that describe how a supplier’s portfolio is allocated across its buyers. The within-link moments are central to the paper’s question: firm-level moments alone do not pin down how much of the scope–buyer co-movement comes from amortization rather than matching (Section 7.1). All results in the main text use this calibration.

The deterministic common-ladder benchmark of Appendix N, in which all buyers purchase prefixes of a common product ranking, serves as a diagnostic. Calibrated to the same data and evaluated on the two estimation economies, it fits far worse than the matching model (criterion 1.45 against 0.36 on the 14 targets both models share). By construction, portfolio coverage is identically one and buyer pairs are never non-nested, against 0.61 and 0.27 in the data, and the supplier’s top product is missing from 3 percent of links against 36 percent. It also cannot generate the scope–buyer elasticity (0.13 against 0.56).

**Calibration method.** Let  $\Theta$  denote the vector of parameters calibrated in the model, and let  $\mathcal{M}$  denote the corresponding set of targeted moments. We estimate  $\Theta$  by simulated method of moments (SMM), minimizing an equally weighted sum of squared relative deviations between simulated and empirical moments:

$$\hat{\Theta} = \arg \min_{\Theta} \sum_{j \in \mathcal{M}} \left( \frac{m_j^{\text{model}}(\Theta) - m_j^{\text{data}}}{m_j^{\text{data}}} \right)^2. \quad (38)$$

Using relative deviations places moments measured in products, buyers, slopes, and shares on a comparable scale. Appendix J.1 describes the distributions of all primitives and the simulation design.

Table 4 summarizes the empirical moments and their model counterparts. It reports the moments used in the SMM objective and additional untargeted moments used to evaluate model fit. The parameters are estimated on two simulated economies; all results reported below evaluate the calibrated parameters on four simulated economies, the two used in the estimation and two additional independent draws. We target seventeen moments. Moments that are not used in the objective function are reported as validation rather than being mechanically fit.

The simulated moments are constructed from realized trade using definitions that mirror those in the data. In particular, a developed product counts toward observed firm scope only if it is sold to at least one buyer, and a buyer counts toward observed buyer base only if positive trade occurs on the supplier–buyer relationship. This distinction matters in the matching model because a developed product may fail to match with any buyer, and different buyers may trade different subsets of the supplier’s portfolio.

**Common empirical targets.** The first group of moments in Table 4 disciplines the distribution of product scope and buyer base. We target the median and the 75th and 90th percentiles of the number of products,  $\hat{G}_s$ , and of the number of buyers,  $N_s$ . In the 2017 manufacturing sample, multiproduct suppliers sell a median of 9 and an average of 16.1 products, while serving a median of 9 and an average of 68.9 buyers. The upper percentiles require the model to reproduce not only the typical supplier but also the strong right skewness in both extensive margins; the means and the standard deviation of scope are reported as validation moments.

We additionally target the median number of buyers per product, defined as the number of realized buyer–product relationships divided by observed product scope. This moment disciplines the joint behavior of buyer reach and product scope rather than either marginal distribution separately. A model could reproduce the average number of buyers and products while generating the wrong way in which the two margins are combined within firms. Buyers per product helps rule out this possibility.

Sales concentration provides a further restriction on the composition of the buyer network. We target the mean buyer-concentration ratio, defined as the share of supplier sales accounted for by its largest buyer. This moment is informative about buyer-quality heterogeneity, networking costs, and selection into the supplier’s active buyer set. We also target the tail index of supplier sales (1.10 in the data), which disciplines the productivity tail  $\theta_\phi$ . In the data it is estimated on the top 1 percent of suppliers; in the simulated economy, which has far fewer firms, on the top 5 percent.

Finally, we target the conditional relationship between buyer base and product scope. Its empirical counterpart is the supplier-fixed-effect specification in Table 2, column (5), in which the coefficient on  $\log N_{st}$  is 0.555. Because the model is static, it has no within-supplier time variation and therefore cannot reproduce the empirical fixed-effect regression literally. We construct its closest model counterpart by estimating, on multiproduct suppliers,

$$\log \hat{G}_s = \beta^{\text{model}} \log N_s + \gamma_1 \log \phi_s + \gamma_2 (\log \phi_s)^2 + \lambda_1 \log c_s + \lambda_2 (\log c_s)^2 + \rho \log \phi_s \log c_s + \varepsilon_s, \quad (39)$$

where the quadratic terms and the interaction allow productivity and the networking cost shifter to affect product scope nonlinearly. Supplier fixed effects in the data absorb all persistent supplier characteristics, whatever their functional form. In the model these characteristics are productivity and the networking cost shifter, so the counterpart controls for both flexibly. The remaining variation in buyer base comes from idiosyncratic buyer-side draws, the static analog of the within-supplier variation that identifies  $\beta$  in the data. The resulting coefficient captures the relationship between buyer base and product scope among suppliers with similar productivity and networking costs. We use this moment to discipline the strength of the buyer–product co-movement, whichever channel generates it.

**Moments specific to product–buyer matching.** The matching model uses these moments to discipline how the supplier’s firm-level product hierarchy maps into individual buyer relationships. The first is the mean product-specialization ratio, defined as the share of supplier sales accounted for by its largest product. The second is within-link rank alignment. Table 11 reports regressions of a product’s rank within a link on its rank in the firm. In the data, the relationship between a product’s importance in the supplier’s overall portfolio and its importance within an individual supplier–buyer link is positive but below one. We target the share-on-share coefficient of 0.82, which is the statistic the model computes. This moment is particularly informative about the dispersion of  $\zeta_{sb}$ . A larger  $\theta_\zeta$  implies less dispersed product–buyer matches and therefore stronger alignment. Together with rank decay, this moment helps distinguish systematic rank decay, governed by  $\alpha$ , from idiosyncratic product–buyer matching, governed by  $\theta_\zeta$ .

**Within-link moments.** Besides alignment, we target five moments that capture how products are distributed across a supplier’s buyer relationships. These are the mean number of products per link (3.24), the share of single-product links (0.42), the decay of within-link sales along the supplier’s lagged product ranking (0.85), the within-supplier dispersion of log link sales (1.52), and portfolio coverage, defined as the share of a supplier’s products sold to its broadest buyer (0.61). The within-link moments are what separate the two channels. Products per link and single-product links pin down how often activation binds, and coverage pins down how differently buyers select from the portfolio. Together they determine how much of the observed scope–buyer co-movement is produced by selection rather than by development.


**Buyer heterogeneity and the calibration of  $D_b$ .** The model allows buyers to differ in profitability through  $D_b$ . In the data, we construct measures of buyer scale independently of the focal supplier, including the buyer’s total purchases and total sales. These variables provide empirical proxies for the expenditure and market-size components summarized by  $D_b$ . Figure 6 shows that buyer heterogeneity matters for supplier sales. The average buyer of a large supplier is only modestly larger than that of a small supplier (panel a), so suppliers of all sizes serve a mix of small and large buyers, but large suppliers’ sales are concentrated on their largest buyers (panel b). This motivates the inclusion of buyer-level profitability heterogeneity in the quantitative model. Appendix A.2 shows that buyers also differ widely in size among the buyers of a single supplier.

We also examine the relationship between these buyer characteristics and supplier–buyer link sales to distinguish persistent buyer-level profitability from pair-specific compatibility, which is captured separately by  $\eta_{sb}$ . In the quantitative implementation, however, we do not map the observed buyer-size measures one-for-one into  $D_b$ . Instead,  $D_b$  is represented by a reduced-form distribution whose dispersion,  $\omega$ , is estimated jointly with the remaining parameters using the firm-level moments in Table 4. The buyer-size evidence in Figure 6 therefore provides empirical motivation and discipline for buyer-level heterogeneity, while the exact dispersion of  $D_b$  is

determined jointly in the SMM calibration.

Consistent with the atomless-supplier assumption in the theoretical model, the buyer-level price index embedded in  $D_b$  is not recomputed in response to the realized shocks of individual suppliers in the finite simulation. The  $S$  simulated suppliers should instead be interpreted as a Monte Carlo approximation to the underlying supplier distribution. Accordingly, we do not renormalize simulated supplier–buyer sales so that they sum exactly to  $E_b$  for each buyer, and buyer-level expenditure exhaustion is not a targeted moment of the current calibration.

Figure 6: Buyer size and supplier size, 2017



*Notes:* Binned scatterplots of the geometric mean of buyers’ total revenue against supplier sales (million USD, log scales), manufacturing suppliers in 2017. Buyer revenue is taken from firms’ administrative tax filings and is therefore independent of the focal supplier’s sales. Panel (a) weights each of a supplier’s buyers equally; panel (b) weights buyers by their purchases from the supplier.

## 5.2 Parameter discipline

Table 3 reports the parameter values, their sources, and the moments that provide the main discipline for each free parameter. Table 4 reports how well the matching model fits both the targeted and untargeted moments. The parameters are jointly identified, since changing any one cost or heterogeneity parameter typically affects several moments at the same time. The moment assignments in Table 3 should therefore be interpreted as the primary sources of discipline rather than as one-to-one identification statements.

The level of the buyer-specific product-activation cost,  $\bar{f}_0^A$ , governs how easily developed products are activated across buyer relationships and is therefore primarily disciplined by the number of products per link and the share of single-product links. The link-cost level,  $\bar{f}^L$ , acts directly on the extensive buyer margin and is primarily disciplined by the buyer-count and buyers-per-product moments. The level of the firm-level product-adoption cost,  $f_0^P$ , determines the overall cost of expanding the supplier’s global portfolio and therefore affects both product scope and, through amortization, the buyer base that can support that scope.

The parameter  $\delta$  governs how the incremental firm-level adoption cost changes as the supplier moves toward more peripheral products. It is therefore particularly important for the dispersion and upper tail of the product-scope distribution. Because  $\delta$  also governs how easily the developed portfolio can expand in response to a larger buyer base, we estimate it rather than fixing it; Appendix L shows that fixing  $\delta = 1.5$  leaves the fit and the main results nearly unchanged. Supplier-level cost heterogeneity, summarized by  $\omega_c$ , generates common variation in networking costs across all buyers of a supplier. It therefore creates variation in buyer base that survives aggregation across links and, through both the matching and the amortization channels, generates co-movement between buyer base and product scope conditional on productivity. For this reason, the within- $\phi$  slope  $\beta$  in equation (39) is an informative moment for  $\omega_c$ . Buyer-quality dispersion,  $\omega$ , instead governs heterogeneity in the profitability of potential buyers and contributes to dispersion in both buyer base and product scope.

In the model,  $\theta_\zeta$  governs product-buyer match dispersion, while  $\alpha$  governs systematic sales decay along the supplier's product ranking. The two parameters are estimated jointly because both affect the relationship between firm-level and link-level product importance. Within-link rank alignment is particularly informative about  $\theta_\zeta$ , while product specialization provides additional discipline on  $\alpha$  and helps separate systematic rank decay from idiosyncratic matching.

The activation-cost rank parameter  $\mu$  is disciplined directly by products per link and the share of single-product links.

The CES elasticity, the productivity floor, and the dispersion of buyer-specific costs are fixed externally or normalized; the productivity tail  $\theta_\phi$  is estimated. Table 3 reports these values and their sources.

**Targeted versus validation moments.** Not every moment reported in Table 4 enters the SMM objective. We report the mean and standard deviation of product scope, mean buyer base, mean buyers per product, the standard deviation of buyer counts, the standard deviation of buyers per product, and the median concentration measures as over-identification checks. These moments are strongly influenced by the upper tail of the supplier distribution and are not assigned to a separate structural parameter in the current calibration.

We also report the share of single-product suppliers as a validation moment. Although the model has no force designed to generate a mass at one product, the model produces a single-product share of 3.8 percent (13.8 in the data), because once selection binds some suppliers find only one product worth developing. We also report the non-prefix share of links, the frequency with which a supplier's top product is missing from a link, and the share of buyer pairs whose product sets are non-nested as untargeted diagnostics of the matching structure. The comparison of targeted moments, untargeted moments, and model-only predictions in Table 4 therefore provides a transparent assessment of both model fit and the dimensions along which additional structure

or empirical discipline may be useful.

Because developed scope is bounded by  $G_{\max}$ , the expected-surplus aggregator is finite for every supplier. For sufficiently productive suppliers every product lies in the region where activation is certain, and expected surplus is linear in  $\phi_s^{\sigma-1}$ . The population moments therefore exist whenever  $\theta_\phi > \sigma - 1 = 3$ . We parameterize the productivity tail as  $\theta_\phi = \theta_\zeta(\sigma - 1) + g$  and estimate  $g > 0$ , so the calibrated tail also satisfies the stronger condition  $\theta_\phi > \theta_\zeta(\sigma - 1)$ , which would be required if all products lay in the interior of the Pareto support. At the calibrated parameters,  $\theta_\phi = 4.63$  against  $\theta_\zeta(\sigma - 1) = 4.07$ .

Table 3: Parameters of the matching model: values, sources, and main sources of discipline

| Parameter                                          | Value                 | Source        | Main sources of discipline                                                                                                               |
|----------------------------------------------------|-----------------------|---------------|------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Fixed or normalized</i>                         |                       |               |                                                                                                                                          |
| $\sigma$ (CES elasticity)                          | 4                     | Literature    | Standard value.                                                                                                                          |
| $\phi_{\min}$ (productivity floor)                 | 1                     | Normalization | Units of productivity.                                                                                                                   |
| $\omega_{fL}, \omega_{fA}$ (buyer-cost dispersion) | 0.6                   | Data          | Residual SD of log buyer degree (0.607) and of buyer mean link scope (0.598).                                                            |
| <i>Calibrated (SMM)</i>                            |                       |               |                                                                                                                                          |
| $\bar{f}_0^A$ (activation cost)                    | 38.4                  | Profiled      | Products per link, one-product links.                                                                                                    |
| $\bar{f}_0^L$ (link cost)                          | 46.3                  | SMM           | Median $N$ , median buyers per product.                                                                                                  |
| $f_0^P$ (adoption cost)                            | $1.45 \times 10^{-5}$ | SMM           | Median and 75th percentile of $\hat{G}_s$ .                                                                                              |
| $\delta$ (adoption convexity)                      | 2.74                  | SMM           | Median, 75th and 90th percentiles of $\hat{G}_s$ .                                                                                       |
| $\mu$ (activation-cost rank slope)                 | -0.487                | SMM           | One-product links, products per link.                                                                                                    |
| $\theta_\zeta$ (match dispersion)                  | 1.36                  | SMM           | Within-link alignment, coverage, specialization.                                                                                         |
| $\alpha$ (rank decay)                              | 0.514                 | SMM           | Rank decay.                                                                                                                              |
| $\omega_\eta$ (pair-cost dispersion)               | 0.648                 | SMM           | Within-supplier SD of log link sales.                                                                                                    |
| $\omega_c$ (networking cost dispersion)            | 0.428                 | SMM           | Median and 75th percentile of $N$ , median buyers per product.                                                                           |
| $\omega$ (buyer-quality dispersion)                | 1.69                  | SMM           | Median $\hat{G}_s$ , 90th percentile of $N$ , buyer concentration.                                                                       |
| $\theta_\phi$ (productivity tail)                  | 4.63                  | SMM           | Sales tail index, 90th percentiles. Estimated through $g$ in $\theta_\phi = \theta_\zeta(\sigma - 1) + g$ , so the tail condition holds. |

*Notes:* The model is calibrated to the 17 moments marked in Table 4, with 10 estimated parameters. The activation-cost level  $\bar{f}_0^A$  is profiled, held at one quarter of its unrestricted estimate, at which the criterion is lower. Identification is joint; the last column lists the most informative moments. Simulation with  $S = 8,000$  suppliers and  $B = 16,000$  potential buyers; moments are averaged over two independent simulated economies. Appendix L reports alternative calibrations.

The calibration fits the typical supplier and the within-link allocation of products closely. The median supplier sells 10.2 products to 11 buyers (9 and 9 in the data), links carry 3.4 products on average (3.2), 37 percent of links carry a single product (42), and portfolio coverage is 0.55 (0.61). Its main misses are on the buyer side and in the heterogeneity of buyers' product sets. The buyer distribution is thinner-tailed than in the data (90th percentile 74 against 100, mean 42 against 69), sales are less concentrated on the largest buyer (0.40 against 0.51) and more evenly spread across a supplier's links (within-supplier standard deviation of log link sales 1.16 against 1.52), and buyers' product sets are non-nested too often (0.58 of buyer pairs against 0.27). On the buyer side, the model reproduces, without targeting it, how link scope rises with buyer size: larger buyers take more products from each supplier, with an elasticity of 0.10 against 0.12 in the data, as the matching channel implies. It overstates, however, how fast the number of a buyer's

Table 4: Identifying moments: data and model

| Moment                                                           | Data     | Model                    |
|------------------------------------------------------------------|----------|--------------------------|
| <i>Product scope</i>                                             |          |                          |
| Median $\hat{G}$                                                 | 9        | 10.2 <sup>†</sup>        |
| Mean $\hat{G}$                                                   | 16.1     | 16.0                     |
| SD of $\hat{G}$                                                  | 21.0     | 18.4                     |
| 75th / 90th percentile of $\hat{G}$                              | 19 / 38  | 18.8 / 34.5 <sup>†</sup> |
| <i>Buyer base</i>                                                |          |                          |
| Median $N$                                                       | 9        | 11 <sup>†</sup>          |
| Mean $N$                                                         | 68.9     | 42.0                     |
| 75th / 90th percentile of $N$                                    | 30 / 100 | 27.5 / 73.5 <sup>†</sup> |
| <i>Link breadth</i>                                              |          |                          |
| Buyers per product (median)                                      | 1.000    | 1.214 <sup>†</sup>       |
| Buyers per product (mean)                                        | 3.749    | 1.447                    |
| <i>Sales composition and product hierarchy</i>                   |          |                          |
| Product specialization ratio (mean)                              | 0.541    | 0.598 <sup>†</sup>       |
| Buyer concentration ratio (mean)                                 | 0.514    | 0.402 <sup>†</sup>       |
| Within-link alignment, link share on firm share                  | 0.820    | 0.655 <sup>†</sup>       |
| <i>Within-link structure</i>                                     |          |                          |
| Products per link (mean)                                         | 3.244    | 3.420 <sup>†</sup>       |
| Share of one-product links                                       | 0.423    | 0.367 <sup>†</sup>       |
| Rank decay (log sales on log lagged rank)                        | 0.852    | 0.682 <sup>†</sup>       |
| Within-supplier SD of log link sales (median)                    | 1.523    | 1.156 <sup>†</sup>       |
| Portfolio coverage $\max_b G_{sb}/\hat{G}_s$ (mean)              | 0.612    | 0.552 <sup>†</sup>       |
| <i>Complementarity and tail</i>                                  |          |                          |
| Scope-buyer elasticity $\beta$ , equation (39) (Table 2, col. 5) | 0.555    | 0.518 <sup>†</sup>       |
| Sales tail index (data top 1%, model top 5%)                     | 1.097    | 0.992 <sup>†</sup>       |
| <i>Untargeted</i>                                                |          |                          |
| SD of $N$                                                        | 784.6    | 174.5                    |
| Buyers per product (SD)                                          | 22.8     | 1.187                    |
| Product specialization ratio (median)                            | 0.512    | 0.579                    |
| Buyer concentration ratio (median)                               | 0.475    | 0.352                    |
| Single-product share                                             | 0.138    | 0.038                    |
| Products per link (SD)                                           | 4.176    | 5.594                    |
| Portfolio coverage (median)                                      | 0.571    | 0.518                    |
| Share of suppliers with full coverage                            | 0.178    | 0.082                    |
| Non-prefix share of links                                        | 0.701    | 0.745                    |
| Top product missing from link (leave-one-out)                    | 0.363    | 0.246                    |
| Non-nested buyer pairs                                           | 0.270    | 0.575                    |
| Link rank score on firm rank score (Table 11, col. 1)            | 0.936    | 1.045                    |
| Residual SD, rank-decay regression                               | 1.227    | 0.744                    |
| $\beta$ also controlling for log sales (col. 6)                  | 0.486    | 0.473                    |
| Products per link on log buyer size (buyer level)                | 0.117    | 0.099                    |
| Number of suppliers on log buyer size (buyer level)              | 0.307    | 0.659                    |
| Targeted moments                                                 |          | 17                       |
| Free parameters                                                  |          | 10                       |
| SMM criterion                                                    |          | 0.466                    |

Notes: <sup>†</sup> targeted. Data moments refer to multiproduct manufacturing suppliers in 2017, except the within-link matching statistics (coverage, non-prefix share, top product missing, alignment, non-nested buyer pairs), which pool 2015–2023. Model moments are computed from realized trade at the calibrated parameters ( $S = 8,000$  suppliers,  $B = 16,000$  potential buyers) and average four simulated economies, two of which were used in the estimation. They refer to multiproduct suppliers except the single-product share and the sales tail index.  $\beta$ : equation (39),  $\log \hat{G}_s$  on  $\log N_s$  with a quadratic in  $\log \phi_s$  and  $\log c_s$  and their interaction. The sales tail index is estimated on the top 1 percent of suppliers in the data and the top 5 percent in the simulated economy, which has far fewer firms. The SMM criterion sums squared relative deviations over the 17 targets. The activation-cost level is profiled and not counted among the free parameters (Table 3).

suppliers rises with buyer size (0.66 against 0.31). Richer buyer-side heterogeneity is a natural extension.

## 6 Quantitative Results

This section presents the quantitative behavior of the calibrated matching model (Section 5). The calibrated values are reported in Table 3. Selection binds in this calibration. Links carry on average 3.4 products, and the broadest buyer takes on average 55 percent of a supplier’s portfolio. Every firm-level moment is computed from realized transactions. Unless noted otherwise, results average four simulated economies at the calibrated parameters. A supplier’s products are the distinct codes it actually sells, and its buyers are those that actually receive at least one product.

The section is organized around the model’s central testable prediction: conditional on productivity, firms with more buyers should have more products. It then inspects the model’s marginal distributions and closes with an untargeted validation exercise based on the growth decompositions of Section 3.

### 6.1 The central prediction: within-productivity co-movement of scope and buyer base

Figure 7 plots the relationship between observed product scope  $\hat{G}_s$  and buyer base  $N_s$  in two ways.

Panel (a) plots  $\log \hat{G}_s$  against  $\log N_s$  separately by productivity quintile. Firms are first divided into five groups based on  $\phi_s$ , and within each group average product scope is plotted across bins of buyer count. The relationship is positive in every productivity quintile. Even among firms with similar productivity, suppliers that reach more buyers also sell more products.

Panel (b) residualizes both variables on the controls of equation (39), a quadratic in  $\log \phi_s$  and  $\log c_s$  and their interaction, the model counterpart of a supplier fixed effect. The residual relationship is clean and positive, with slope  $\beta = 0.518$ . Table 5 traces the same object. The unconditional elasticity is 0.641 (column 1). Controlling for productivity barely changes it, whether linearly (0.636, column 2) or with twenty productivity bins (0.631, column 3). The elasticity falls only when the networking cost is also controlled for: to 0.601 with linear controls (column 4) and to 0.518 with the flexible controls of equation (39) (column 5), against 0.555 in the data. The decline reflects that suppliers with high networking costs have both fewer buyers and fewer products, a common driver that supplier fixed effects absorb in the data.


The variance decomposition in Appendix K Table 13 sharpens the point. Productivity alone explains 59% of the variance of  $\log \hat{G}_s$ , and of the remaining within- $\phi$  variation, buyer count absorbs about 76%. Table 14 shows the within-quintile correlation between  $\log \hat{G}_s$  and  $\log N_s$

rising from 0.82 in the lowest productivity quintile to 0.97 in the highest.

The within-productivity slope is the model’s central prediction, but by itself it does not reveal which channel generates it. A firm with more buyers may develop more products because the adoption cost is spread over more relationships (amortization), or may sell more of an unchanged portfolio because each product has more chances to find a buyer (matching). Both produce the same slope. Section 7.1 separates them. The unconditional slope alone would be uninformative, because more productive firms mechanically have both more buyers and more products; the conditional slope instead isolates buyer–product co-movement from productivity and networking cost sorting. By itself, however, it does not distinguish amortization from matching.

**Match with the data.** The model’s 0.518 compares with an empirical within-supplier elasticity of 0.555 (Table 2, column 5). Controlling in addition for log sales, the model gives 0.473 against 0.486 in the data.<sup>7</sup>

Figure 7: Within-productivity co-movement of product scope and buyer base



*Notes:* Calibrated matching model, four simulated economies, multiproduct suppliers. Panel (a) plots  $\log \hat{G}_s$  against  $\log N_s$  within productivity quintiles. Panel (b) residualizes both variables, within each economy, on the controls of equation (39), a quadratic in  $\log \phi_s$  and  $\log c_s$  and their interaction; the fitted slope is  $\beta = 0.518$ , as in Table 5, column (5).

<sup>7</sup>The coefficient in column (5) of Table 5 is 0.475, 0.544, 0.523 and 0.531 in the four simulated economies. Table 5 and the text report their average. This variation across simulated economies is larger than the within-economy standard errors in Table 5. With twenty productivity bins and linear  $\log c_s$  instead of the quadratic controls, the average is 0.529.

Table 5: Regressions of  $\log \hat{G}_s$  on  $\log N_s$ ,  $\log \phi_s$ , and  $\log c_s$  (calibrated model)

|                                         | (1)              | (2)              | (3)              | (4)               | (5)              |
|-----------------------------------------|------------------|------------------|------------------|-------------------|------------------|
| $\log N_s$                              | 0.641<br>(0.003) | 0.636<br>(0.004) | 0.631<br>(0.004) | 0.601<br>(0.010)  | 0.518<br>(0.010) |
| $\log \phi_s$                           |                  | 0.035<br>(0.025) |                  | 0.214<br>(0.052)  |                  |
| $\log c_s$                              |                  |                  |                  | -0.077<br>(0.019) |                  |
| $\phi$ -bin FE                          | No               | No               | Yes              | No                | No               |
| Quadratic in $\log \phi_s$ , $\log c_s$ | No               | No               | No               | No                | Yes              |
| $R^2$                                   | 0.897            | 0.897            | 0.900            | 0.897             | 0.904            |

*Notes:* Dependent variable  $\log \hat{G}_s$  (products sold to at least one buyer), multiproduct suppliers.  $\phi$ -bin FE are twenty  $\log \phi_s$  bins. Column (5) is equation (39):  $\log \phi_s$ ,  $\log c_s$ , their squares and their interaction, the calibrated counterpart of Table 2, column (5). Coefficients, standard errors and  $R^2$  average four simulated economies at the calibrated parameters. Standard errors in parentheses are within-economy and do not reflect simulation variation across economies, which is larger (footnote 7).

## 6.2 Marginal distributions and calibration fit

Figure 8 reports the simulated distributions of firm-level product scope  $\hat{G}_s$ , buyer count  $N_s$ , and total sales  $T_s$ . Product scope and buyer base are both strongly right-skewed, with long upper tails, while the distribution of log sales is approximately bell-shaped.


The model reproduces the central tendency of the two extensive margins well. Among multiproduct suppliers, the median simulated supplier sells 10.2 products and serves 11 buyers, compared with 9 products and 9 buyers in the 2017 data. Thus, the model slightly overstates median product scope and median buyer base in the calibration sample. The dispersion and upper percentiles of product scope are also close to the data (standard deviation 18.4 against 21.0; 90th percentile 34.5 against 38), and the model produces a single-product share of 3.8 percent against 13.8 percent. The buyer distribution is thinner-tailed than in the data (90th percentile 74 against 100). We view these discrepancies as limitations of the quantitative fit rather than as objects central to the buyer-product complementarity mechanism.

Appendix K.2 shows that simulated sales co-move positively with buyer base, product scope, and supplier productivity, consistent with the corresponding cross-sectional patterns in the data.

## 6.3 Untargeted validation: the composition of growth

The calibration targets firm-level distributional moments in 2017 together with within-link allocation statistics and the scope-buyer association measured in the transaction panel. The growth decompositions of 3 are not targeted and therefore provide an untargeted validation exercise for the model's mechanics. Although the model is static, it admits a natural comparative-static analog.

Figure 8: Marginal distributions of firm-level outcomes in the calibrated model



*Notes:* The figure reports simulated distributions of firm-level product scope  $\hat{G}_s$  (panel a), buyer count  $N_s$  (panel b), and log total sales  $\log T_s$  (panel c) in the calibrated model, for all active suppliers in one simulated economy, including single-product suppliers. Dashed lines indicate medians (10 products and 10 buyers). Table 4 reports moments for multiproduct suppliers only, to match the empirical sample, which is why its medians are higher. For display, panel (a) omits suppliers with more than 150 products and panel (b) suppliers with more than 300 buyers.


We raise every supplier’s productivity by 10 percent, re-solve each supplier’s buyer and product decisions, and decompose the resulting change in sales into the same four buyer  $\times$  product margins used in the data, taking the supplier’s baseline buyer and product sets as the “old” reference. Because the shock weakly expands every supplier’s buyer and product sets, the decomposition is exact.

Figure 9 shows that the model reproduces the main qualitative gradient in the data. Expansion through the extensive margins is concentrated among the fastest-growing, least-productive and initially smallest suppliers. In the top growth percentile, new buyers purchasing products already in the portfolio account for about 93 percent of growth (panel a). The joint addition of new buyers and new products is most important among small suppliers, where it accounts for about 16 percent of growth (panel c), and declines with size and productivity. By contrast, the largest and most productive suppliers grow mostly along the intensive margin: existing buyer–product pairs account for about 72 percent of growth in the top size tercile and 73 percent in the top productivity quintile. As in the data (Figure 4), the composition of expansion therefore shifts from new buyers and products toward established buyer–product relationships as suppliers become larger.

The fit is qualitative rather than quantitative. The joint new-buyer–new-product margin is considerably smaller in the model than in the data, consistent with the model’s limited product-development response (Section 7.1). Two further caveats apply. First, the static model has no product or link exit, so the bottom of the simulated growth distribution contains suppliers that expand least rather than suppliers that contract. Second, because the shock is common and positive, the informative object is the composition of each supplier’s adjustment rather than the distribution of growth rates. We therefore read the figure as an untargeted validation of how

the model’s adjustment varies across suppliers, not as a structural reproduction of the 2017–2023 transition.

Figure 9: Composition of growth after a +10% productivity shock (calibrated model)



*Notes:* Comparative-static analog of the empirical growth decomposition in Section 3. Every supplier’s productivity  $\phi_s$  is raised by 10 percent and the supplier’s buyer and product decisions are re-solved. The resulting change in sales is decomposed into new-link-new-product (NLNP), new-link-old-product (NLOP), old-link-new-product (OLNP), and old-link-old-product (OLOP) sales, where “old” refers to the supplier’s baseline buyer and product sets. Within each bin, components are summed across suppliers and divided by the bin’s total sales growth, so the four shares sum to one. Panel (a) groups suppliers by their simulated growth rate, panel (b) by baseline productivity quintile, and panel (c) by baseline size tercile. The decomposition is untargeted.

## 7 Counterfactuals

### 7.1 Decomposing the scope–buyer elasticity: amortization versus matching

The within-productivity co-movement of product scope and buyer base can arise through two channels (Section 4.5). Under amortization, a supplier with more buyers develops more products because the firm-level adoption cost is spread over more relationships. Under matching, a supplier with more buyers sells more of a given portfolio because each product has more chances to find a buyer. Both deliver the same regression coefficient. The model lets us separate them.

**The fixed-portfolio counterfactual.** We hold each supplier’s developed portfolio fixed at the portfolio that suppliers with its productivity and networking costs develop. The portfolio keeps its size but can no longer respond to the supplier’s own buyers, so amortization is shut down, while matching operates as before. We then re-estimate  $\beta$  on the resulting observed scope. The fall in  $\beta$  measures amortization; the remainder measures matching.

**Results.** Table 6 reports the decomposition. The elasticity falls only from 0.518 to 0.513, so amortization accounts for about 1 percent of the co-movement (standard error 2.2 percentage

points across simulated economies) and matching for the rest. Suppliers with more buyers do develop more products: the elasticity of developed scope with respect to buyers is 21 percent of  $\beta$ . But each new product is less appealing than the previous one, so the products added in response to more buyers are peripheral and rarely sold, and the response barely reaches observed scope.

**Extension: product appeal.** Because new products are always less appealing than existing ones, the benchmark is a lower bound on amortization. Appendix [M](#) relaxes this by letting each product carry its own appeal, common to all of the supplier’s buyers, so that a newly developed product can be more attractive than earlier ones. Calibrations of this extension that fit the data comparably well imply an amortization share between 6 and 12 percent of  $\beta$ , depending on how much products differ in appeal, which the targeted moments do not determine precisely. Matching therefore accounts for the large majority of the co-movement in every specification.

**Robustness.** Within the benchmark, the result does not depend on the curvature of the adoption cost. With  $\delta$  fixed at 1.5 instead of the estimated 2.74, the developed-portfolio response rises to 29 percent of  $\beta$ , but amortization still accounts for only 4.5 percent of observed co-movement, with essentially the same fit (Appendix [L](#)).

**What identifies the decomposition.** Firm-level moments, including the scope–buyer elasticity itself, identify the decomposition only weakly: both channels imply that suppliers with more buyers sell more products. Calibrations that target only firm-level moments can settle in configurations in which the elasticity is reproduced with weak selection, so that links carry many products, almost no link carries a single product, the broadest buyer takes nearly the whole portfolio, and amortization accounts for most of the co-movement. These configurations fit the firm-level moments less well than ones close to our calibration, but the firm-level moments alone do not rule them out sharply. The within-link moments measure how tightly selection binds and do.

**Evidence from incumbent buyers.** The incumbent-buyer exercise provides an additional diagnostic. Products sold to a supplier’s incumbent buyers cannot be added by matching with new buyers, so their response to buyer base reflects product development. In the data, the incumbent elasticity is 0.38, about 69 percent of the total (Appendix [D](#)). Development therefore plays a larger role in the data than the static model generates. A static model in which products are developed once and matched to buyers ex post can reproduce the within-link allocation of products only by loading the co-movement onto selection. The incumbent-buyer evidence suggests that suppliers respond to a larger buyer base by developing and extending products within existing relationships over time, a force that a dynamic model of product adoption would capture.

Table 6: Sources of the scope–buyer elasticity: amortization versus matching

|                                                                        | Model      |
|------------------------------------------------------------------------|------------|
| $\beta$ , calibrated economy                                           | 0.518      |
| $\beta$ , developed portfolio fixed                                    | 0.513      |
| Amortization share of $\beta$                                          | 1.0% (2.2) |
| Developed-portfolio response ( $\beta$ on $G_s^*$ , share of $\beta$ ) | 21%        |
| Data: share of co-movement among incumbent buyers                      | 69%        |

*Notes:* The fixed-portfolio counterfactual sets each supplier’s developed portfolio to the portfolio that suppliers with its productivity and networking costs develop and re-estimates  $\beta$  (equation (39)) on observed scope. The amortization share is the fall in  $\beta$  relative to the calibrated economy; the standard error in parentheses, in percentage points, reflects variation across simulated economies. The developed-portfolio response is the coefficient of the same regression with developed scope  $G_s^*$  as the dependent variable, relative to  $\beta$ . When products also differ in appeal, the amortization share is between 6 and 12 percent (Appendix M, Table 18). The data row is the coefficient of log products sold to incumbent buyers on log buyers (0.382, Table 9, column 5) relative to the total elasticity (0.555, Table 2, column 5); it is an upper bound on amortization, since development can respond to buyer base for other reasons.

**Removing the adoption cost altogether.** Setting  $F^P = 0$  is not a useful test of the mechanism, because without a firm-level cost suppliers develop products up to the enumeration ceiling. Observed scope nevertheless stays bounded, because selection determines which developed products are sold, and the elasticity rises to 0.59. This is the matching channel at work: even when every product has been developed, suppliers with more buyers sell more of them. The common-ladder benchmark behaves differently (0.13 calibrated, 0.11 with  $F^P = 0$ ), because its scope is pinned by the deepest buyer’s cutoff (Appendix N).

## 7.2 Comparative statics in the fixed costs

Having decomposed the channels that tie the two margins together, we now ask how the calibrated economy responds to reductions in the fixed costs governing each margin.

Table 7 reports the response to halving, in turn, the firm-level product-adoption cost  $F^P$  and the buyer-relationship cost  $\bar{f}^L$ , holding all other parameters fixed.

Spillovers run mainly from buyers to products. Halving the buyer-relationship cost raises buyer base by about 45 percent and product scope by about 14 percent. Halving the product-adoption cost raises product scope by about 7 percent but buyer base by only about 2 percent. Because the link between the margins runs mainly through matching, a larger buyer base raises observed scope primarily by helping existing products find buyers rather than by inducing substantial new development; in the other direction, a broader portfolio adds little to the value of marginal buyer relationships.

These comparative statics carry a policy reading that depends on the decomposition. A policy that lowers buyer–supplier matching costs raises not only the buyer base but also the number

of products suppliers sell. In the calibration disciplined by within-link moments, however, the additional products mostly reflect existing products finding new buyers rather than new product development. Had the co-movement been driven by amortization, the same policy would stimulate product development. Distinguishing the two channels is therefore not only a question of interpretation, it determines whether easing network frictions is also an innovation policy. This is the quantitative counterpart of the policy argument in the introduction, and it is a channel that a single-margin model, in which product scope and buyer acquisition are chosen independently, cannot deliver. Because the model is partial equilibrium, these responses abstract from adjustments in buyers' price indices.

We vary only  $F^P$  and  $\bar{f}^L$  because they govern the two extensive margins directly. The third fixed cost, the per-product activation cost  $\bar{f}^A$ , plays a different role; it determines how developed products are allocated across buyer relationships. In the calibration disciplined by within-link moments, this cost is large enough that activation binds on most links, which is what makes product–buyer selection, and hence the matching channel, quantitatively important. Its level is pinned down by the number of products per link and the share of single-product links, rather than by the firm-level margins studied here.

Table 7: Comparative statics in the two fixed costs

| Cost reduction                                | Change relative to baseline |                  |
|-----------------------------------------------|-----------------------------|------------------|
|                                               | Product scope $\hat{G}_s$   | buyer base $N_s$ |
| Product-adoption cost $F^P$ : $-50\%$         | $+7\%$                      | $+2\%$           |
| Buyer-relationship cost $\bar{f}^L$ : $-50\%$ | $+14\%$                     | $+45\%$          |

*Notes:* Each row re-solves the calibrated economy after reducing one fixed cost by 50 percent, holding all other parameters fixed. Entries report percentage changes in average product scope and buyer base relative to the baseline calibrated economy.

**Shutting down one extensive margin at a time.** The comparative statics above move the two margins jointly. A complementary exercise shuts each margin down entirely and asks how much of the dispersion in firm sales depends on it. With the buyer margin shut down, every supplier sells its portfolio to a single representative buyer of average quality, so cross-firm sales variation comes only from productivity and product scope. With the product margin shut down, every supplier sells a single product to its buyers, so variation comes only from productivity and the buyer base. Both are computed on the calibrated economy, holding each firm's realized objects fixed, so productivity remains active in both and neither distribution collapses to a point.

Table 8 reports the result, and it is asymmetric. Shutting down the buyer margin cuts the standard deviation of log sales by 47 percent and the top 1 percent sales share by two thirds; shutting down the product margin cuts dispersion by only 11 percent. The two counterfactuals

are not a decomposition, since productivity operates through both margins and their effects need not sum to one, but they show that dispersion in buyer bases accounts for most of the dispersion in firm size, with product scope amplifying rather than originating it. The last column gives every buyer the same quality while keeping the realized buyer and product sets. Dispersion then rises by 20 percent: in the model, the marginal buyers of suppliers with many buyers are small, and giving each of them average quality makes large suppliers larger still. What drives firm-size dispersion is therefore the number of buyers a supplier reaches, not the quality of the buyers it serves. The model does, however, understate the role of the intensive margin: sales per buyer–product pair account for about half of the variance of log supplier sales in the data but for about 16 percent in the model, consistent with its too-even spread of sales across a supplier’s links (Table 4).

Table 8: Shutting down one extensive margin at a time

|                                | Full model | Buyer margin<br>shut down | Product margin<br>shut down | Equal buyer<br>quality |
|--------------------------------|------------|---------------------------|-----------------------------|------------------------|
| SD(log $T_s$ )                 | 1.709      | 0.897                     | 1.513                       | 2.051                  |
| change relative to full model  |            | −47.5%                    | −11.4%                      | +20.1%                 |
| share of full-model dispersion | 100%       | 52.5%                     | 88.6%                       |                        |
| Top 1% sales share             | 0.638      | 0.215                     | 0.569                       | 0.901                  |
| change relative to full model  |            | −66.2%                    | −10.9%                      | +41.3%                 |

*Notes:* Accounting exercise on the calibrated economy, holding each supplier’s realized buyer set, product set and link-level product sets fixed, with match quality set to its mean. Averages over four simulated economies; the ranking of the columns is the same in every economy. Full model:  $\phi_s^{\sigma-1} \sum_b D_b \sum_{g \in \mathcal{G}_{sb}} h(g)^{1-\sigma}$ . Buyer margin shut down (every supplier sells its portfolio to one representative buyer of mean quality):  $\phi_s^{\sigma-1} \bar{D} \sum_{g \in \hat{\mathcal{G}}_s} h(g)^{1-\sigma}$ , where  $\hat{\mathcal{G}}_s \equiv \bigcup_{b \in \mathcal{B}_s} \mathcal{G}_{sb}$ . Product margin shut down (every supplier sells a single product to its buyers):  $\phi_s^{\sigma-1} \sum_{b \in \mathcal{B}_s} D_b$ . Equal buyer quality: the full-model expression with  $D_b = \bar{D}$  for every buyer. The standard deviation of log realized sales, with actual match draws, is 1.711.

## 8 Conclusion

This paper studies the joint role of buyer relationships and product scope in shaping firm size inside production networks. Using the universe of firm-to-firm transactions in Chile, we show that these are empirically distinct but closely connected margins. Buyer relationships have shorter and less persistent life cycles than supplier–product relationships, and large supplier expansions disproportionately combine new buyers with new products. More importantly for the model, buyer base and product scope co-move strongly within suppliers. Conditional on supplier fixed effects, a 10 percent increase in the number of buyers is associated with a 5.6 percent increase in product scope, about two thirds of it among products sold to incumbent buyers. The same pattern appears for product entry. Suppliers tend to introduce more products in years when they also serve more buyers. The response is concentrated among peripheral products, consistent with buyer expansion raising the return to products near the margin of adoption.

We develop a static model in which suppliers jointly choose their firm-level product portfolio and the buyers they serve, and in which buyer base raises product scope through two channels. Under amortization, a firm-level product-adoption cost paid once generates returns across many buyer relationships. Under matching, buyer-specific product–buyer matches give each developed product more chances to find a buyer as the buyer base grows. The data support the need for buyer-specific matching. Buyers largely agree on which of a supplier’s products are best, but a supplier’s broadest buyer purchases only part of its portfolio, its top product is missing from many of its links, and different buyers of the same supplier often purchase different, non-nested subsets of products.

The model reproduces the within-supplier relationship between buyer base and product scope and allows us to decompose it into its two underlying channels. The decomposition depends on which features of the data discipline the model. Firm-level moments alone identify the two channels only weakly. Once the model is also disciplined by how products are allocated across buyer relationships, matching accounts for the large majority of it: amortization accounts for about 1 percent in the benchmark and between 6 and 12 percent in an extension in which products differ in appeal. The empirical decomposition among incumbent buyers, however, points to a larger role for product development than the static model generates. Because this response occurs within existing buyer relationships, it cannot be explained by the mechanical matching of existing products to new buyers. This gap suggests that buyer expansion also affects firms’ product-development decisions through forces that are not fully captured by the static model. These forces may reflect a stronger dynamic amortization mechanism or other forms of adjustment in product development. Distinguishing among them requires a dynamic framework in which buyer relationships and product portfolios evolve jointly over time.

The quantitative exercises also show that the interaction between the two margins matters for the distribution of firm sales. Spillovers run mainly from buyers to products. Lowering buyer-relationship costs raises product scope substantially, whereas lowering product-adoption costs barely affects buyer networks. The two margins are not equally important quantitatively. Removing the buyer margin reduces the dispersion of log sales by 47 percent, compared with 11 percent when the product margin is removed. Equivalently, the buyer margin alone reproduces about 89 percent of the sales dispersion generated by the full model, while the product margin alone reproduces about 53 percent. buyer base is therefore the main source of firm-size dispersion in the calibrated economy, with product scope amplifying differences generated through the buyer network.

Although some of our empirical facts are measured over time, the benchmark model is static. We therefore interpret its growth exercises as comparative statics rather than as a model of firm dynamics. Nevertheless, when the calibrated economy is perturbed, the model reproduces the qualitative reallocation across old and new buyers and products observed in the data: smaller and

less productive suppliers adjust relatively more through extensive margins, whereas larger and more productive suppliers rely more heavily on existing buyer–product relationships. A natural extension is therefore to introduce buyer accumulation and product adjustment dynamically. Such a model would allow us not only to speak directly to the life-cycle patterns documented in the data, but also to explain the incumbent-buyer evidence and determine whether the additional product-development response reflects dynamic amortization or other forces linking the evolution of customer relationships to product adoption.

More broadly, the results show that product scope cannot be studied separately from the network of buyers over which products are sold. Policies that reduce frictions to forming supplier–buyer relationships raise the number of products suppliers sell; whether they also stimulate product development depends on the relative strength of amortization and matching, which within-link and incumbent-buyer evidence can discipline. Policies that lower the cost of product adoption, by contrast, have little effect on buyer networks in the calibrated economy. Accounting for these interactions is therefore important for understanding how buyer–supplier networks and multiproduct production jointly shape the distribution of firm activity.

## References

- Aekka, Anuraag and Gaurav Khanna (2024) “Endogenous Production Networks and Firm Dynamics,” Technical report, JMP.
- Afrouzi, Hassan, Andres Drenik, and Ryan Kim (2023) “Concentration, Market Power, and Misallocation: The Role of Endogenous Customer Acquisition,” Working Paper 31415, National Bureau of Economic Research, [10.3386/w31415](#).
- Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P. Vasquez (2022) “The Effects of Joining Multi-national Supply Chains: New Evidence from Firm-to-Firm Linkages,” *The Quarterly Journal of Economics*, 137 (3), 1495–1552, [10.1093/qje/qjac006](#).
- Argente, David, Munseob Lee, and Sara Moreira (2018) “How do Firms Grow? The Life Cycle of Products Matters,” 2018 Meeting Papers 1174, Society for Economic Dynamics, <https://EconPapers.repec.org/RePEc:red:sed018:1174>.
- Arkolakis, Costas, Sharat Ganapati, and Marc-Andreas Muendler (2021) “The Extensive Margin of Exporting Products: A Firm-Level Analysis,” *American Economic Journal: Macroeconomics*, 13 (4), 182–245, [10.1257/mac.20150370](#).
- Arkolakis, Costas, Federico Huneus, and Yuhei Miyauchi (2023a) “Production Network Formation, Trade, and Welfare,” NBER Working Paper 30954, National Bureau of Economic Research.

- (2023b) “Spatial Production Networks,” Working Paper 30954, National Bureau of Economic Research, [10.3386/w30954](#).
- 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., Andreas Moxnes, and Karen Helene Ulltveit-Moe (2018) “Two-Sided Heterogeneity and Trade,” *Review of Economics and Statistics*, 100 (3), 424–439.
- Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott (2010) “Multiple-Product Firms and Product Switching,” *American Economic Review*, 100 (1), 70–97, [10.1257/aer.100.1.70](#).
- (2011) “Multiproduct Firms and Trade Liberalization,” *Quarterly Journal of Economics*, 126 (3), 1271–1318.
- Burstein, Ariel, Javier Cravino, and Marco Rojas (2024) “Input price dispersion across buyers and misallocation\*,” *Working Paper*.
- Cai, Jing and Adam Szeidl (2018) “Interfirm Relationships and Business Performance,” *The Quarterly Journal of Economics*, 133 (3), 1229–1282, [10.1093/qje/qjx049](#).
- Cavenaile, Laurent and Pau Roldan-Blanco (2021) “Advertising, Innovation, and Economic Growth,” *American Economic Journal: Macroeconomics*, 13 (3), 251–303, [10.1257/mac.20180461](#).
- Copeland, Adam and Adam Hale Shapiro (2016) “Price Setting and Rapid Technology Adoption: The Case of the PC Industry,” *The Review of Economics and Statistics*, 98 (3), 601–616, [10.1162/REST\\_a.00539](#).
- Criscuolo, Chiara, Ralf Martin, Henry G. Overman, and John Van Reenen (2019) “Some Causal Effects of an Industrial Policy,” *American Economic Review*, 109 (1), 48–85, [10.1257/aer.20160034](#).
- Demir, Banu, Ana Cecília Fielér, Daniel Yi Xu, and Kelly Kaili Yang (2024) “O-Ring Production Networks,” *Journal of Political Economy*, 132 (1), 200–247.
- Eaton, Jonathan, Samuel S. Kortum, and Francis Kramarz (2026) “Firm-to-Firm Trade: Imports, Exports, and the Labor Market,” *Econometrica*, 94 (4), 1135–1170.
- Einav, Liran, Peter J. Klenow, Jonathan D. Levin, and Raviv Murciano-Goroff (2021) “Customers and Retail Growth,” NBER Working Papers 29561, National Bureau of Economic Research, Inc, <https://ideas.repec.org/p/nbr/nberwo/29561.html>.

- Hopenhayn, Hugo A. (1992) “Entry, Exit, and firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60 (5), 1127–1150, <http://www.jstor.org/stable/2951541>.
- Hottman, Colin J., Stephen J. Redding, and David E. Weinstein (2016) “Quantifying the Sources of Firm Heterogeneity,” *The Quarterly Journal of Economics*, 131 (3), 1291–1364, <https://doi.org/10.1093/qje/qjw012>.
- Huneus, Federico (2018a) “Production Network dynamics and the propagation of shocks,” *Working Paper*.
- (2018b) “Production Network Dynamics and the Propagation of Shocks,” Working paper.
- Juhász, Réka, Nathan Lane, and Dani Rodrik (2024) “The New Economics of Industrial Policy,” *Annual Review of Economics*, 16, 213–242, [10.1146/annurev-economics-081023-024638](https://doi.org/10.1146/annurev-economics-081023-024638).
- 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, [10.1093/qje/qjz024](https://doi.org/10.1093/qje/qjz024).
- Mayer, Thierry, Marc J. Melitz, and Gianmarco I. P. Ottaviano (2014) “Market Size, Competition, and the Product Mix of Exporters,” *American Economic Review*, 104 (2), 495–536.
- Miyauchi, Yuhei (2024) “Matching and Agglomeration: Theory and Evidence from Japanese Firm-to-Firm Trade,” *Econometrica*, 92 (6), 1869–1905.

# Appendix

## A Buyer Breadth Relative to Product Scope

Figure 10 provides a direct counterpart to the buyer–product density measures reported in Figure 2. For each supplier–year, let  $B_{st}$  denote the number of buyers,  $G_{st}$  the number of products, and  $R_{st}$  the number of active buyer–product pairs.

The average number of buyers per product is  $R_{st}/G_{st}$ , while the average number of products per buyer is  $R_{st}/B_{st}$ . Their ratio is

$$\frac{R_{st}/G_{st}}{R_{st}/B_{st}} = \frac{B_{st}}{G_{st}}.$$

Thus,  $B_{st}/G_{st}$  is not itself the average number of buyers per product. Instead, it measures buyer base relative to product scope: how large the supplier’s customer base is compared with the size of its product portfolio.

Figure 10 shows that this ratio increases systematically with supplier sales. Larger suppliers therefore expand their customer base more rapidly than their product portfolio. This provides a direct summary of the asymmetry documented in Figure 2, where the average number of buyers reached by each product rises more strongly with supplier size than the average number of products sold within each buyer relationship.

### A.1 Buyer and Product Margins in Supplier-Size Dispersion

This appendix provides a complementary cross-sectional decomposition of supplier-size dispersion into buyer and product margins. The exercise asks how much of the variation in supplier sales is associated with differences in the number of relationships or products, as opposed to differences in average sales conditional on those extensive margins.

On the buyer side, observed supplier sales can be written as

$$S_s = B_s \times \bar{S}_s^b,$$

where  $B_s$  is the number of buyers served by supplier  $s$  and  $\bar{S}_s^b = S_s/B_s$  is average sales per buyer. In logs,

$$\log S_s = \log B_s + \log \bar{S}_s^b.$$

We use this identity to decompose the variance of log supplier sales into a component associated with the number of buyers and a component associated with average sales per buyer.

Analogously, on the product side,




Figure 10: Buyer Breadth Relative to Product Scope

*Notes:* The figure presents a binned scatterplot relating supplier size to the ratio of the number of buyers to the number of products,  $B_{st}/G_{st}$ . Supplier size is measured by sales. Both axes are shown on a log scale. Each dot represents a bin of suppliers based on sales, and the solid line reports the fitted relationship.

$$S_s = G_s \times \bar{S}_s^g,$$

where  $G_s$  is the number of products sold and  $\bar{S}_s^g = S_s/G_s$  is average sales per product. Hence,


$$\log S_s = \log G_s + \log \bar{S}_s^g,$$

which gives the corresponding decomposition into product scope and average sales per product.

Figure 11 reports both decompositions for the 2017 cross section. The buyer extensive margin accounts for 42.9 percent of the variance of log supplier sales, while average sales per buyer accounts for the remaining 57.1 percent. In the analogous product decomposition, product scope accounts for 27.3 percent of firm-size dispersion and average sales per product for 72.7 percent. Thus, both extensive margins are quantitatively important, but differences in buyer base account for a substantially larger share of cross-sectional supplier-size dispersion than differences in product scope.

These decompositions should be interpreted separately. The buyer and product shares are alternative ways of decomposing the same variation in supplier sales and therefore should not be added across decompositions. The comparison is instead informative about the relative importance of customer reach and product scope for cross-sectional differences in supplier size.

Figure 11: Buyer and product margins in supplier-size dispersion



*Notes:* The figure reports two alternative decompositions of the variance of log supplier sales in the 2017 cross section. The buyer decomposition separates the contribution of the number of buyers from average sales per buyer. The product decomposition separates the contribution of the number of products from average sales per product. The number of buyers accounts for 42.9 percent of the variance of log supplier sales and sales per buyer for 57.1 percent. The corresponding product-side contributions are 27.3 percent and 72.7 percent, respectively. Each bar represents a separate decomposition and sums to 100 percent; contributions should not be added across the buyer and product decompositions.


## A.2 Within-supplier heterogeneity in buyer size

Figure 12 complements Figure 6. For each supplier, it computes the standard deviation of log buyer size across the supplier’s buyers. The within-supplier standard deviation of log buyer size is typically between 2 and 3, indicating substantial size heterogeneity among buyers of the same supplier. This motivates modeling buyer profitability  $D_b$  as heterogeneous across the buyers a supplier serves.

## B Dynamics of Buyer Links and Supplier–Product Relationships

This appendix provides additional evidence that supplier–buyer links and supplier–product relationships are empirically distinct objects. Buyer links turn over more quickly and exhibit shorter sales lifecycles than supplier–product relationships. These patterns motivate the model’s separation of buyer base and product scope. Because the benchmark model is static, however, the

Figure 12: Within-supplier dispersion in buyer size, 2017



*Notes:* Distribution across manufacturing suppliers of the standard deviation, across each supplier’s active buyers, of log buyer input spending (left) and log buyer total revenue (right) in 2017. Buyer input spending excludes purchases from the focal supplier, so the measure does not mechanically reflect the size of the supplier–buyer relationship. Buyer total revenue is obtained from firms’ administrative tax filings. Suppliers with at least five buyers and nonmissing buyer-size information. Each observation is one supplier; the vertical axes report density.

lifecycle profiles and survival patterns documented here are not quantitative targets of the model.

## B.1 Relationship age and cohort measurement

For the lifecycle analysis, we define the age of a supplier–buyer link as the number of years since that pair is first observed trading:

$$\text{Age}_{sbt}^L = t - \min \{ \tau : S_{sbt}^{\text{link}} > 0 \} .$$

Analogously, the age of a supplier–product relationship is

$$\text{Age}_{sgt}^P = t - \min \{ \tau : S_{sgt}^{\text{prod}} > 0 \} .$$

Supplier age is defined separately as the number of years since the supplier is first observed with positive sales:

$$\text{Age}_{st}^F = t - \min \{ \tau : S_{st}^{B2B} > 0 \} .$$

Because the transaction panel begins in 2015, relationship entry is observed without left-censoring only for links and supplier–product relationships that begin after the start of the sample. Lifecycle regressions therefore use relationships whose first observed activity occurs within the sample and should be interpreted as describing observed relationship age rather than the complete

age of relationships that may have existed before 2015.

Buyer links and supplier–product relationships follow different lifecycle patterns. For supplier–buyer links, we estimate

$$\log S_{sbt}^{\text{link}} = \alpha^L + \sum_{j \geq 1} \beta_j^L \mathbf{1}\{\text{Age}_{sbt}^L = j\} + \mu_{sb} + \gamma_t + \varepsilon_{sbt}^L, \quad (40)$$


where age zero is the omitted category,  $\mu_{sb}$  is a supplier–buyer-link fixed effect, and  $\gamma_t$  is a year fixed effect. For supplier–product relationships, we estimate

$$\log S_{sgt}^{\text{prod}} = \alpha^P + \sum_{j \geq 1} \beta_j^P \mathbf{1}\{\text{Age}_{sgt}^P = j\} + \mu_{sg} + \gamma_t + \varepsilon_{sgt}^P, \quad (41)$$

where  $\mu_{sg}$  is a supplier–product fixed effect. Standard errors are clustered at the relationship level in each specification.


Figure 13 plots the estimated age coefficients. Both relationships display hump-shaped sales profiles, but buyer links peak earlier and decline more quickly. Link sales peak at approximately 0.13 log points above their entry-year level at age one. Supplier–product sales continue rising through age two, reach a peak of approximately 0.20 log points, and remain above the link profile at subsequent ages. By age five, the supplier–product premium remains close to 0.09 log points, whereas the link premium has fallen to approximately 0.03.

Figure 13: Sales over the lifecycle of buyer links and supplier–product relationships




*Notes:* The figure plots the estimated coefficients from equations (40) and (41). The link specification includes supplier–buyer-link and year fixed effects; the product specification includes supplier–product and year fixed effects. Age zero is the omitted category, so each coefficient measures the average within-relationship change in log sales relative to the first observed year of the relationship. Shaded bands report 95 percent confidence intervals. Standard errors are clustered at the corresponding relationship level.


## B.2 Evolution of total sales by cohort and year




(a) Link total sales per cohort and year



(b) Product total sales per cohort and year



(a) Link survival rate



(b) Product survival rate

### B.3 Survival rate of links and products

## C Additional Decompositions of Supplier Growth

### C.1 Composition of Extensive-Margin Growth

The main decomposition in Section 3.1 includes the OLOP component, which combines changes in sales within the supplier’s baseline buyer–product space with losses from exiting relationships. To focus more directly on adjustment involving new buyers or new products, this appendix reports the three extensive-margin components—NLOP, OLNP, and NLNP—without plotting OLOP.

For each supplier, define the total extensive-margin growth as

$$\Delta S_s^{\text{EXT}} = \Delta S_s^{\text{NLOP}} + \Delta S_s^{\text{OLNP}} + \Delta S_s^{\text{NLNP}}.$$

Then, for suppliers with nonzero extensive-margin adjustment, define the contribution of margin  $c$  to total extensive-margin sales growth as

$$\tilde{s}_s^c = \frac{\Delta S_s^c}{\Delta S_s^{\text{EXT}}}, \quad c \in \{\text{NLOP}, \text{OLNP}, \text{NLNP}\}.$$


Because each component is normalized by total extensive-margin adjustment,  $\Delta S_s^{\text{EXT}}$ , the shares need not all be positive. A negative contribution indicates that a particular extensive margin moves in the opposite direction from the supplier’s net extensive-margin adjustment.

Figure 16 reports these contributions across the same intervals of the realized-growth distribution used in Figure 3. The composition of extensive-margin adjustment changes markedly across the growth distribution. In the lower tail, some extensive-margin components contribute negatively to the net adjustment. As realized supplier growth increases, however, the composition shifts increasingly toward NLNP. Among the fastest-growing suppliers, a large share of extensive-margin adjustment is associated with transactions involving both buyers and products that were

not present in the baseline year. NLOP also remains quantitatively important, whereas OLNP contributes relatively little.

The figure therefore reinforces the main Fact 1 result from a different perspective. Even after focusing only on adjustment involving new buyers or new products, the joint new-buyer-new-product margin becomes especially important among the largest expansions. The fastest-growing suppliers do not simply expand through buyer acquisition or product introduction separately; their extensive-margin growth disproportionately combines both margins.

Figure 16: Composition of Extensive-Margin Growth




*Notes:* The figure reports the shares of positive extensive-margin growth accounted for by NLNP, NLOP, and OLNP across intervals of the realized supplier-growth distribution. The groups are defined using adjacent percentile cutoffs, as in Figure 3.

## C.2 The Role of Relationship Exit

A further decomposition of OLOP distinguishes changes in sales among old buyer-product relationships that remain active through 2023 from losses associated with relationships that disappear between 2017 and 2023. Figure 17 shows that relationship exit is the dominant source of contraction in the lower tail of the realized-growth distribution. Among the suppliers with the largest declines, nearly all of the fall in sales is accounted for by the disappearance of baseline buyer-product relationships, while changes in sales among continuing OLOP relationships play only a small role. In some of the other declining groups, losses from destroyed relationships exceed the supplier's net contraction because they are partially offset by sales generated through new buyers or new products.

The pattern changes markedly among growing suppliers. Losses from relationship exit continue to offset some expansion, but become progressively smaller relative to total growth. At the same time, growth increasingly comes from new buyer–product combinations, particularly NLNP, while continuing OLOP relationships also contribute positively. Thus, the asymmetry across the growth distribution reflects not only differences in the importance of extensive-margin expansion, but also substantial differences in relationship destruction: contraction is driven primarily by the disappearance of existing buyer–product relationships, whereas the largest expansions rely increasingly on the formation of new ones.

Figure 17: Decomposition of OLOP Growth into Continuing and Destroyed Relationships



*Notes:* The figure decomposes supplier sales growth between 2017 and 2023 across intervals of the realized-growth distribution. The old link–old product (OLOP) component is separated into changes in sales among baseline buyer–product relationships that remain active in 2023 (continuing) and losses from baseline buyer–product relationships that are no longer active in 2023 (destroyed). The remaining components are new link–new product (NLNP), new link–old product (NLOP), and old link–new product (OLNP). Each component is normalized by the supplier’s total change in sales. For shrinking suppliers, losses from destroyed relationships therefore enter with a positive share because both the loss and total sales growth are negative, while positive sales generated through new buyers or products appear as negative shares because they offset the overall contraction. B denotes intervals in the bottom half and T intervals in the top half of the realized-growth distribution; categories are defined using adjacent percentile cutoffs as in Figure 3.

### C.3 Heterogeneity by Initial Supplier Sales

We next examine whether the composition of supplier growth varies with suppliers’ initial scale. We divide suppliers into low-, middle-, and high-sales categories based on their sales in 2017 and repeat the growth decomposition across the four buyer–product margins. Figure 18 reports


the results. Panel (a) shows the full decomposition, while Panel (b) excludes OLOP and focuses on the composition of growth across the three extensive margins.

The composition of growth varies systematically with initial supplier size. Among initially smaller suppliers, NLNP accounts for the largest share of growth, so expansion disproportionately combines new buyers with new products. As initial sales increase, the importance of NLNP declines, while growth through existing buyer-product relationships becomes more important. Larger suppliers therefore rely relatively less on the joint expansion of buyer and product scope and more on margins involving products already in their portfolios.


Panel (b) shows that this pattern is not driven solely by the greater importance of OLOP among larger suppliers. Restricting attention to the extensive margins, the contribution of NLNP declines with initial sales, while NLOP becomes increasingly important. Thus, smaller suppliers are relatively more likely to expand buyer and product scope jointly, whereas larger suppliers are more likely to reach new buyers with products they already sell. OLNP remains comparatively small across all three sales categories.

This size gradient provides a useful target for the model’s heterogeneous comparative statics. In particular, the model should generate a shift from the joint expansion of buyer and product scope among smaller suppliers toward greater reliance on established products as supplier scale increases.

Figure 18: Growth decomposition by sales category



(a) Growth decomposition by sales category



(b) Growth decomposition by sales without OLOP

## C.4 Heterogeneity over the Supplier Life Cycle


We then ask whether the composition of growth also varies over the supplier life cycle, while conditioning on suppliers’ sales in 2017. This exercise separates differences associated with supplier age from those arising simply because older suppliers tend to operate at a different initial scale. Figure 19 reports the decomposition for young, middle-aged, and old suppliers.

The clearest age pattern appears in the composition of extensive-margin growth. Younger suppliers rely relatively more on NLNP, jointly adding new buyers and new products, whereas the importance of NLOP increases with supplier age. Older suppliers are therefore relatively more likely to expand by forming new buyer relationships for products already included in their existing portfolios. OLNP remains small across age groups.


Panel (b) makes this shift particularly clear. After excluding OLOP, the share of extensive-margin growth accounted for by NLNP falls monotonically with supplier age, while the share accounted for by NLOP rises. Conditional on initial sales, younger suppliers therefore expand buyer and product scope more jointly, whereas older suppliers increasingly expand their customer base using products they have already adopted.

These patterns are consistent with a life-cycle interpretation in which younger suppliers are still jointly developing their customer base and product portfolio, while more mature suppliers increasingly extend the market reach of an established product portfolio. Because supplier age is not an explicit state variable in our static model, however, we interpret this exercise as complementary descriptive evidence rather than as a moment that the model is required to match directly.

Figure 19: Growth decomposition by age



(a) Growth decomposition by age



(b) Growth decomposition by age without OLOP

*Notes:* The figure decomposes supplier sales growth between 2017 and 2023 into NLNP, NLOP, OLNP, and OLOP by supplier age group, conditioning on suppliers' sales in 2017. Young suppliers are between one and five years old, middle-aged suppliers are between six and fifteen years old, and old suppliers are more than fifteen years old. Supplier age is measured as of 2017.

Taken together, the initial-size and age decompositions show that the relative importance of buyer and product expansion varies systematically across suppliers. Initially smaller suppliers expand relatively more through the joint addition of new buyers and new products, whereas larger suppliers rely more heavily on established products and existing buyer-product relationships. A related pattern appears over the supplier life cycle: conditional on initial sales, younger suppliers

expand relatively more through the joint addition of buyers and products, while older suppliers increasingly extend established products to new customers. The initial-size results provide a direct target for the model’s heterogeneous comparative statics, while the age patterns provide complementary descriptive evidence on how these margins vary with supplier maturity.

## D Buyer breadth and product scope among incumbent buyers

A potential concern with the relationship documented in Section 3.3 is that the positive association between buyer base and product scope could arise mechanically from the formation of new buyer relationships. If newly acquired buyers purchase products that the supplier does not sell to its existing customers, an increase in the number of buyers would mechanically coincide with an increase in measured product scope, even in the absence of broader adjustments across the supplier’s existing relationships.

To assess this possibility, we repeat the analysis measuring product scope using only products sold to incumbent buyers. An incumbent buyer is a buyer with an existing relationship with the supplier prior to year  $t$ . The dependent variable is therefore the log number of distinct products sold by supplier  $s$  in year  $t$  to its incumbent buyers. We estimate specifications analogous to those in Table 2, progressively including sector and supplier fixed effects and controlling for contemporaneous supplier sales.

Table 9 shows that the relationship remains positive and statistically significant when product scope is restricted to incumbent relationships. In the specification with supplier and year fixed effects, the elasticity of product scope to incumbent buyers with respect to buyer base is 0.382. Controlling additionally for contemporaneous supplier sales reduces the coefficient to 0.288. Thus, conditional on persistent supplier heterogeneity and current firm size, a 10 percent increase in buyer base is associated with approximately a 2.9 percent increase in the number of products sold to incumbent buyers.

The magnitude is smaller than for total product scope, where the corresponding coefficient in the specification that also controls for contemporaneous supplier sales is 0.486 (Table 2, column 6). This difference is consistent with part of the overall co-movement between buyer base and product scope occurring through products associated with newly formed buyer relationships. Importantly, however, the positive relationship remains when these newly formed relationships are excluded from the outcome. Suppliers that serve more buyers also tend to operate a broader product portfolio within their pre-existing customer relationships.

This exercise therefore suggests that the Fact 3 relationship is not driven solely by the mechanical addition of products accompanying new buyers. Instead, changes in buyer base are associated with broader product scope even among established relationships. As in the baseline analysis, we

interpret this evidence as a conditional association rather than a causal effect of buyer acquisition. The result reinforces the motivation for a model in which buyer and product scope are jointly determined and adjustments along one margin can affect incentives along the other.

Table 9: Buyer Breadth and Product Scope Among Incumbent Buyers

| <i>Dependent variable: Log Number of Products Sold to Incumbent Buyers</i> |                       |                        |                       |                       |                       |                       |
|----------------------------------------------------------------------------|-----------------------|------------------------|-----------------------|-----------------------|-----------------------|-----------------------|
|                                                                            | (1)                   | (2)                    | (3)                   | (4)                   | (5)                   | (6)                   |
| Log Number of Buyers                                                       | 0.447***<br>(0.00190) | 0.402***<br>(0.00225)  | 0.463***<br>(0.00179) | 0.387***<br>(0.00219) | 0.382***<br>(0.00459) | 0.288***<br>(0.00487) |
| Log Sales                                                                  |                       | 0.0692***<br>(0.00169) |                       | 0.116***<br>(0.00179) |                       | 0.185***<br>(0.00388) |
| Observations                                                               | 106,300               | 106,143                | 106,299               | 106,142               | 104,166               | 103,983               |
| Firms                                                                      | 21,627                | 21,622                 | 21,626                | 21,621                | 19,493                | 19,462                |
| $R^2$                                                                      | 0.330                 | 0.340                  | 0.452                 | 0.474                 | 0.862                 | 0.867                 |
| Year FE                                                                    | Yes                   | Yes                    | Yes                   | Yes                   | Yes                   | Yes                   |
| Sector FE                                                                  | No                    | No                     | Yes                   | Yes                   | No                    | No                    |
| Supplier FE                                                                | No                    | No                     | No                    | No                    | Yes                   | Yes                   |

*Notes:* The dependent variable is the log number of distinct products (CPCs) sold by supplier  $s$  in year  $t$  to incumbent buyers. The main explanatory variable is the log number of distinct buyers served by the supplier in the same year. Columns (1)–(2) include year fixed effects; columns (3)–(4) additionally include supplier-sector fixed effects; and columns (5)–(6) include supplier fixed effects instead. Even-numbered columns additionally control for contemporaneous log supplier sales. Heteroskedasticity-robust standard errors are reported in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

## E Buyer Breadth and Product Additions

In the main text, Fact 3 documents a strong positive relationship between a supplier’s buyer base and its overall product scope. This appendix shows that the same pattern is also present at the margin of product additions. We define a product as new to a supplier in year  $t$  if the supplier sells that CPC in year  $t$  but did not sell it in year  $t - 1$ . Thus, this measure captures additions to the supplier’s product portfolio relative to the previous year rather than the product’s first-ever appearance in the data.

Table 10 relates the log number of product additions to the supplier’s contemporaneous buyer base. The specifications mirror those in Table 2 in the main text. Columns (1)–(2) include year fixed effects, columns (3)–(4) additionally control for supplier-sector fixed effects, and columns (5)–(6) instead include supplier fixed effects. Even-numbered columns additionally control for contemporaneous supplier sales.

The relationship between buyer base and product additions is positive and precisely estimated across all specifications. Most importantly, it remains strong after introducing supplier fixed

effects, which absorb persistent differences across firms. In the specification with year and supplier fixed effects, the coefficient on log buyers is 0.392. Controlling additionally for contemporaneous sales reduces the coefficient to 0.340, but it remains highly statistically significant. Therefore, within the same supplier and conditional on current firm size, a 10 percent increase in buyer base is associated with approximately a 3.4 percent increase in the number of product additions.

These results complement the evidence in the main text. The positive relationship between buyer base and product scope is not driven solely by differences in firms' existing portfolios: suppliers with broader customer bases also add more products. As in the main analysis, we interpret these estimates as conditional associations rather than causal effects of buyer acquisition on product adoption. They provide additional evidence for the model-facing fact that buyer base and product scope systematically co-move within suppliers.

Table 10: Buyer Breadth and Product Additions

| <i>Dependent variable: Log Number of New Products</i> |                       |                        |                       |                        |                       |                        |
|-------------------------------------------------------|-----------------------|------------------------|-----------------------|------------------------|-----------------------|------------------------|
|                                                       | (1)                   | (2)                    | (3)                   | (4)                    | (5)                   | (6)                    |
| Log Buyers                                            | 0.261***<br>(0.00149) | 0.246***<br>(0.00171)  | 0.289***<br>(0.00136) | 0.255***<br>(0.00163)  | 0.392***<br>(0.00331) | 0.340***<br>(0.00362)  |
| Log Sales                                             |                       | 0.0217***<br>(0.00129) |                       | 0.0491***<br>(0.00128) |                       | 0.0948***<br>(0.00271) |
| Observations                                          | 138,784               | 138,534                | 138,784               | 138,534                | 138,219               | 137,966                |
| Firms                                                 | 23,382                | 23,380                 | 23,382                | 23,380                 | 22,817                | 22,812                 |
| $R^2$                                                 | 0.225                 | 0.226                  | 0.420                 | 0.426                  | 0.735                 | 0.739                  |
| Year FE                                               | Yes                   | Yes                    | Yes                   | Yes                    | Yes                   | Yes                    |
| Sector FE                                             | No                    | No                     | Yes                   | Yes                    | No                    | No                     |
| Supplier FE                                           | No                    | No                     | No                    | No                     | Yes                   | Yes                    |

*Notes:* The dependent variable is the log number of products sold by supplier  $s$  in year  $t$  that were not sold by the supplier in year  $t - 1$ . The main explanatory variable is the log number of distinct buyers served by the supplier in year  $t$ . Columns (1)–(2) include year fixed effects; columns (3)–(4) additionally include supplier-sector fixed effects; and columns (5)–(6) include supplier fixed effects instead. Even-numbered columns additionally control for contemporaneous log supplier sales. Heteroskedasticity-robust standard errors are reported in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.10$ .

## F Evidence on Cannibalization


The two dimensions of firm scope studied in the paper raise two distinct cannibalization concerns. First, when a supplier introduces an additional product to an existing buyer, sales of the new product may come at the expense of products that the supplier was already selling to that buyer. Second, when a supplier acquires a new buyer, sales to the new relationship may come at

the expense of its incumbent buyer relationships. This appendix examines these two possibilities separately.

## F.1 Product Introductions within Existing Buyer Relationships

We first ask whether introducing an additional product within an existing supplier–buyer relationship displaces sales of the supplier’s incumbent products. Figure 20 provides a descriptive view of the event. We define  $t = 0$  as the first year in which the supplier adds a new product within an existing buyer relationship and follow sales of the products already sold in that relationship and of the newly introduced products. Restricting attention to supplier–buyer links that survive throughout the event window, incumbent-product sales do not decline following the introduction. Instead, sales of old and new products expand together.

Figure 20: Sales around a product introduction within an existing link




*Notes:* The figure plots mean log sales to the incumbent buyer for old products and new products in years surrounding the first addition of a new product within an existing supplier–buyer relationship. Event time  $t = 0$  denotes the product introduction. Old products are those sold within the relationship before the new product is introduced; new products are those first added at  $t = 0$ . The sample contains supplier–buyer relationships that remain active throughout the event window. The figure is descriptive and plots raw sales paths rather than estimates relative to a comparison group.

Figure 21 complements these raw sales paths with a balanced-pre comparison. For each product-introduction cohort, we define incumbent products as those with positive sales in every pre-event year from  $t = -4$  through  $t = -1$ . Treated relationships are compared with never-treated and not-yet-treated relationships, using  $t = -1$  as the common baseline. To make sales changes comparable across relationships, product and total-link sales are normalized by total sales in the supplier–buyer relationship at  $t = -1$ . We focus here on relationships that remain active through the end of the sample, so that changes in incumbent products are not mechanically generated by the disappearance of the supplier–buyer relationship itself.

Panel [a](#) shows a sharp and persistent expansion in total supplier–buyer sales following the introduction of a new product. Sales of balanced-pre incumbent products change little on impact but decline at longer horizons. The expansion in total relationship sales therefore exceeds the displacement of incumbent-product sales, although the latter is not absent.

Panel [b](#) shows that the fraction of incumbent products remaining active also declines at longer horizons. Together, the two panels indicate gradual portfolio turnover: new products expand the relationship while some incumbent products lose sales or cease to be sold. The evidence supports limited immediate cannibalization, but not the absence of longer-run displacement.

Figure 21: Product introduction and incumbent products: balanced-pre comparisons



*Notes:* The figure studies the first addition of a new product within an existing supplier–buyer relationship. Incumbent products are products with positive sales in each of the four years preceding the event,  $t = -4, \dots, -1$ . The sample shown here is restricted to supplier–buyer relationships that remain active through the end of the sample. For each treatment cohort and event horizon, treated relationships are compared with never-treated and not-yet-treated relationships, with  $t = -1$  as the common baseline. Panel (a) plots estimated effects on sales of balanced-pre incumbent products and on total supplier–buyer sales, with both outcomes normalized by total link sales at  $t = -1$ . Panel (b) reports the estimated effect on the fraction of balanced-pre incumbent products that continue to have positive sales. Vertical bars report 95 percent confidence intervals; standard errors are clustered at the supplier level.

Taken together, Figures [20](#) and [21](#) show that new-product sales are associated with a substantial expansion of the supplier–buyer relationship. This expansion is not fully offset by losses on incumbent products. Nevertheless, Figure 21 indicates that incumbent-product sales and retention decline at longer horizons. We therefore interpret the evidence as showing limited immediate displacement alongside gradual portfolio turnover. These patterns motivate the separable benchmark and the nested-CES extension of Appendix [O](#) as an approximation; they do not establish that within-link cannibalization is absent or that the within-supplier and across-supplier elasticities are equal.


## F.2 New-Buyer Introductions and Incumbent Buyer Relationships

The second exercise studies cannibalization along the buyer margin. We ask whether acquiring a new buyer diverts activity away from the supplier’s existing buyer relationships. Here, the event is different from the exercise above:  $t = 0$  denotes the arrival of a new buyer rather than the introduction of a new product within an existing relationship. Incumbent links are supplier–buyer relationships that were already active before the new buyer arrives.

Figure 22 examines two potential forms of displacement. Panel **a** asks whether the arrival of a new buyer makes incumbent relationships more likely to disappear. We find no systematic decline in incumbent-link survival. If anything, the estimates are positive after the event, although they are imprecisely estimated.

Panel **b** instead considers the intensive allocation of the supplier’s sales across buyers. The sales share accounted for by incumbent buyers falls on impact when the new buyer arrives, as one would mechanically expect when an additional buyer enters the supplier’s customer portfolio. The decline is temporary: the incumbent share subsequently recovers and is close to its pre-event level within roughly two years. Thus, the addition of a new buyer generates some short-run reallocation of the supplier’s sales across customers, but we find no evidence of persistent displacement of incumbent buyer relationships.

Figure 22: New-buyer introduction and incumbent buyer relationships



*Notes:* The figure examines incumbent buyer relationships around the arrival of a new buyer. Event time  $t = 0$  denotes the year in which the new buyer enters the supplier’s customer portfolio. Incumbent links are buyer relationships that were active before the event. Panel (a) reports the estimated effect on the probability that an incumbent relationship remains active in each post-event year. Panel (b) reports the evolution of incumbent buyers’ share of supplier sales around the event, with  $t = -1$  as the omitted reference period. The sample is restricted to cohorts observed at every event time shown in each panel. Vertical bars report 95 percent confidence intervals; standard errors are clustered at the supplier level.

The two exercises address different forms of potential displacement. Following product introduction within an existing relationship, total relationship sales expand substantially, while

incumbent-product sales and retention decline at longer horizons. Following new-buyer acquisition, we find no systematic decline in incumbent-link survival. Incumbent buyers temporarily account for a smaller share of supplier sales, but their share subsequently recovers.

These results suggest that expansion is not merely a redistribution of a fixed amount of activity across products or customers. They nevertheless allow for gradual displacement within product portfolios. Because the buyer exercise measures survival and sales shares, it does not by itself establish the absence of displacement in incumbent buyers' sales levels.

## G Rank structure

The matching structure has a testable implication for how a supplier's product hierarchy maps into individual buyer relationships. If buyers' tastes were purely idiosyncratic, a product's rank within a link would be unrelated to its rank in the supplier's overall sales. Under a common product ladder, by contrast, every buyer purchases a prefix of the same ranking, so the supplier's top product would be sold on every link and lead almost all of them. Table 11 shows that the data lie in between. Within a link, a product's rank is strongly related to its rank in the firm (column 1: a slope of 0.86–0.94), so buyers largely agree on which of a supplier's products are best. But the agreement is far from complete: the supplier's top product is only 49 percentage points more likely than its other products to be a given buyer's top product from that supplier (column 3). Consistently, the top product is missing from 36 percent of the supplier's links, and 27 percent of pairs of buyers of the same supplier purchase non-nested product sets (Table 4). This combination of a common ranking with buyer-specific deviations is what the matching model generates: rank decay  $\alpha$  governs the common ranking, while the dispersion of the match shock  $\theta_\zeta$  governs the deviations. The calibrated model reproduces the high agreement without targeting it, with a score slope of 1.05 against 0.94 in the data.

Appendix H shows that the buyer–product complementarity documented in Table 2 is not uniform across the supplier's product portfolio. We study product-level sales around the addition of a large new buyer and allow the response to vary with the product's pre-event position in the supplier's portfolio. Products are classified as core, middle, or peripheral according to their pre-event sales rank, with core products serving as the omitted category. The response is strongly concentrated toward the periphery of the portfolio. Before the arrival of the new buyer, sales of middle and peripheral products evolve similarly to those of core products. Beginning in the event year, however, peripheral products experience a sharp and persistent increase in sales relative to core products. The response of middle products is positive but substantially smaller. The effect therefore becomes progressively stronger as we move from core toward peripheral products.

This heterogeneity is consistent with both channels in the model (Section 4.5). Under matching, a new buyer gives products near the margin of the portfolio, which are sold to few buyers, a new

Table 11: Within-link product-rank alignment

|                                     | (1)                 | (2)                 | (3)                 | (4)                 | (5)                 | (6)                 |
|-------------------------------------|---------------------|---------------------|---------------------|---------------------|---------------------|---------------------|
| Dependent variable                  | Link score          | Link rank           | Link top-1          | Link top-1          | Link top-3          | Link top-3          |
| Regressor                           | Firm score          | Firm rank           | Firm top-1          | Firm score          | Firm top-3          | Firm score          |
| <i>Panel A: Link FE and year FE</i> |                     |                     |                     |                     |                     |                     |
| Coefficient                         | 0.863***<br>(0.031) | 0.144***<br>(0.005) | 0.491***<br>(0.032) | 0.589***<br>(0.024) | 0.409***<br>(0.025) | 0.830***<br>(0.057) |
| <i>Panel B: Link×year FE</i>        |                     |                     |                     |                     |                     |                     |
| Coefficient                         | 0.936***<br>(0.038) | 0.139***<br>(0.005) | 0.492***<br>(0.035) | 0.606***<br>(0.028) | 0.398***<br>(0.028) | 0.829***<br>(0.065) |
| Observations                        | 56,565,645          | 56,565,645          | 56,565,645          | 56,565,645          | 48,832,461          | 48,832,461          |
| Suppliers                           | 52,334              | 52,334              | 52,334              | 52,334              | 45,023              | 45,023              |

*Notes:* The unit of observation is a supplier–buyer–product–year. For each supplier, products are ranked by their total sales across all buyers (firm rank) and by their sales within each supplier–buyer link (link rank). Scores convert ranks to a scale from 1 (top) to 0 (bottom):  $\text{score} = 1 - (\text{rank} - 1) / (\text{number of products} - 1)$ , where the firm score is computed over all of the supplier’s products and the link score over the products sold on that link. Top-1 (top-3) dummies equal one for the top (top three) products. Because a link’s score spans the full  $[0, 1]$  range while its products cover only part of the firm ranking, the score slope can exceed one even if every buyer follows the firm ranking. Standard errors clustered by supplier in parentheses. \*\*\*  $p < 0.01$ .

opportunity to be sold. Under amortization, a new buyer raises the return to maintaining and developing those products. The evidence therefore complements the product-introduction results in Table 10: buyer acquisition is associated disproportionately with activity in peripheral products rather than with further expansion of the products that already account for the largest share of the supplier’s sales.

## H Buyer Acquisition and Relative Sales across the Product Ladder

This appendix examines how the sales of a supplier’s existing products evolve around the acquisition of a large new buyer. We ask whether peripheral products experience stronger or weaker sales growth than core and middle products. This distinguishes changes concentrated in previously low-selling products from changes shared across the supplier’s portfolio. The exercise concerns relative sales performance among continuing products; product disappearance is a separate outcome that these samples do not measure.

We define the event as the first observed year in which a supplier adds a large buyer with whom it has not previously traded in the observed data. Buyer size is measured by total purchases in the preceding year, and large buyers are those at or above the annual 75th percentile of the buyer-size distribution. Sales are aggregated across all buyers within each supplier–product–year, so the outcome captures a product’s total B2B sales rather than its sales exclusively to the newly acquired buyer.

We consider two samples. The event-window sample retains products with strictly positive sales in every year from  $t = -4$  through  $t = 4$ . The full-sample survival specification imposes the stronger requirement of positive sales in every calendar year of the dataset. Within each supplier’s eligible product basket, products are ranked by their average sales over the available pre-event years. Core products comprise approximately the top 10 percent, middle products the next 40 percent, and peripheral products the remaining 50 percent, subject to integer rounding. These groups remain fixed after the event. Because the eligible basket differs across samples, a product’s rank category can also differ between specifications.

The dependent variable is log supplier–product sales. Regressions include supplier–product and supplier–year fixed effects, and standard errors are clustered at the supplier level. Core products and event time  $t = -1$  are the reference categories. A positive coefficient for peripheral products therefore indicates that their sales increased more, or decreased less, than core-product sales relative to the year before buyer acquisition. It does not establish an absolute increase in peripheral sales or an absolute decline in core sales.


Figure 23 reports the event-window results. Peripheral-product estimates are positive in most post-event years, but fluctuate and their confidence intervals generally include zero. Middle-product estimates are also imprecise. The figure therefore provides suggestive evidence of stronger relative sales performance for peripheral products, but does not establish a persistent or precisely estimated ordering across the product ladder. The uncertainty in the pre-event estimates also limits the conclusions that can be drawn about differential pre-trends.

Figure 24 reports results for products active throughout the full sample. Peripheral-product sales exhibit a larger, persistent increase relative to core-product sales after buyer acquisition,

while the relative increase for middle products is smaller. However, some pre-event coefficients also differ from zero at the plotted 95 percent confidence level. Moreover, the sample selects products using their continued activity throughout the dataset. The stronger response therefore describes this selected product basket and should not be interpreted as a causal effect of the new buyer or a result that is invariant to the product-survival restriction.


The figures show that the clearest relative gains for peripheral products occur among products that remain active throughout the full sample. This is consistent with buyer acquisition being associated with changes in the relative commercial importance of existing products, although the event-window evidence is weaker. Neither figure establishes whether peripheral products become more or less likely to disappear: products that cease to be sold during the required period are excluded by construction. Likewise, the relative coefficients cannot determine whether peripheral gains occur alongside expansion or contraction in core-product sales. Assessing product survival or cannibalization requires retaining the initial product basket when sales fall to zero and tracking its sales and activity directly.

Figure 23: Relative Product Sales around Buyer Acquisition: Event-Window Sample



*Notes:* The figure reports differential changes in log supplier–product B2B sales around the first observed addition of a new relationship with a large buyer. Large buyers are defined using lagged total purchases and an annual 75th-percentile cutoff. Products have positive sales in every year from  $t = -4$  through  $t = 4$ . Within each supplier’s eligible product basket, Core products are approximately the top 10 percent by average pre-event sales, Middle products the next 40 percent, and Peripheral products the remaining 50 percent. Groups are fixed before the event. Core products and  $t = -1$  are the reference categories. Regressions include supplier–product and supplier–year fixed effects. Vertical bars report 95% confidence intervals based on standard errors clustered at the supplier level.

Figure 24: Relative Product Sales around Buyer Acquisition: Products Active throughout the Full Sample



*Notes:* The figure uses the same outcome, buyer-size definition, fixed effects, and reference categories as Figure 23, but retains only products with positive sales in every calendar year of the full dataset. Rank groups are constructed using average pre-event sales within this eligible basket and remain fixed after the event. The sample and product ranks can therefore differ from those in Figure 23. Vertical bars report 95% confidence intervals based on standard errors clustered at the supplier level.

## H.1 CES aggregator and expenditure minimization

For each active supplier–product variety  $(s, g)$ , let  $q_{sbg}$  denote the quantity purchased by buyer  $b$  and  $p_{sbg}$  its price. Buyer–product match quality is denoted by  $\zeta_{sbg} > 0$ . Buyer  $b$  combines the available supplier–product varieties according to

$$M_b = \left[ \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg}^{1/\sigma} q_{sbg}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1. \quad (42)$$

The normalization  $\zeta_{sbg}^{1/\sigma}$  implies that the corresponding quantity demand and revenue equations are linear in the product–buyer match shifter  $\zeta_{sbg}$ .

To derive the associated price index, consider the expenditure-minimization problem for a target level  $\bar{M}_b$ :

$$\min_{\{q_{sbg}\}} \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} p_{sbg} q_{sbg} \quad (43)$$

subject to

$$\left[ \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg}^{1/\sigma} q_{sbg}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \geq \bar{M}_b. \quad (44)$$

Let  $\lambda_b$  denote the multiplier on the CES constraint. The Lagrangian is

$$\begin{aligned} \mathcal{L}_b = & \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} p_{sbg} q_{sbg} \\ & + \lambda_b \left\{ \bar{M}_b - \left[ \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg}^{1/\sigma} q_{sbg}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \right\}. \end{aligned} \quad (45)$$

For any active variety  $(s, g)$ , the first-order condition is

$$p_{sbg} = \lambda_b \bar{M}_b^{1/\sigma} \zeta_{sbg}^{1/\sigma} q_{sbg}^{-1/\sigma}. \quad (46)$$

Rearranging,

$$q_{sbg} = \bar{M}_b \zeta_{sbg} \left( \frac{\lambda_b}{p_{sbg}} \right)^\sigma. \quad (47)$$

Substituting (47) into (44) and imposing equality at the optimum gives

$$1 = \lambda_b^{\sigma-1} \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg} p_{sbg}^{1-\sigma}. \quad (48)$$

Hence

$$\lambda_b = \left[ \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg} p_{sbg}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (49)$$

The multiplier is therefore the buyer-level CES price index,

$$P_b = \left[ \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg} p_{sbg}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (50)$$

Equivalently,

$$P_b^{1-\sigma} = \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg} p_{sbg}^{1-\sigma}. \quad (51)$$

Using  $\lambda_b = P_b$  in (47), Hicksian demand for an active supplier–product variety is

$$q_{sbg}^H = \bar{M}_b \zeta_{sbg} \left( \frac{p_{sbg}}{P_b} \right)^{-\sigma}. \quad (52)$$

The associated expenditure function is

$$\sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} p_{sbg} q_{sbg}^H = P_b \bar{M}_b. \quad (53)$$

Thus  $P_b$  is the unit-expenditure function for buyer  $b$ 's intermediate-input bundle.

## H.2 Quantity demand and product-level revenue

In the benchmark partial-equilibrium environment, buyer  $b$  has total expenditure  $E_b$ . Expenditure exhaustion implies

$$E_b = P_b M_b, \quad \text{or equivalently} \quad M_b = \frac{E_b}{P_b}. \quad (54)$$

Substituting (54) into (52) gives quantity demand:

$$\begin{aligned} q_{sbg} &= \frac{E_b}{P_b} \zeta_{sbg} \left( \frac{p_{sbg}}{P_b} \right)^{-\sigma} \\ &= E_b \zeta_{sbg} p_{sbg}^{-\sigma} P_b^{\sigma-1}. \end{aligned} \quad (55)$$

Multiplying quantity demand by the product price gives expenditure on supplier–product variety  $(s, g)$ , which is also the supplier's product-level revenue from buyer  $b$ :

$$\begin{aligned}
r_{sbg} &\equiv p_{sbg} q_{sbg} = E_b \zeta_{sbg} p_{sbg}^{1-\sigma} P_b^{\sigma-1} \\
&= E_b \zeta_{sbg} \left( \frac{p_{sbg}}{P_b} \right)^{1-\sigma}.
\end{aligned} \tag{56}$$

Equation (56) is the product-level revenue equation used in the main text. The corresponding share of buyer  $b$ 's expenditure allocated to product  $(s, g)$  is

$$\frac{r_{sbg}}{E_b} = \zeta_{sbg} \left( \frac{p_{sbg}}{P_b} \right)^{1-\sigma}. \tag{57}$$

Summing expenditure shares across all varieties available to buyer  $b$  gives

$$\begin{aligned}
\sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \frac{r_{sbg}}{E_b} &= P_b^{\sigma-1} \sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} \zeta_{sbg} p_{sbg}^{1-\sigma} \\
&= P_b^{\sigma-1} P_b^{1-\sigma} \\
&= 1,
\end{aligned} \tag{58}$$

where the second equality uses (51). Thus total expenditure is exhausted across the active supplier-product varieties.

### H.3 Pricing and the supplier-side revenue equation

For completeness, this appendix also shows how (56) maps into the reduced-form revenue equation used in the supplier's problem.

Product  $g$  of supplier  $s$  has productivity

$$\phi_{s,g} = \frac{\phi_s}{h(g)}, \tag{59}$$

so its marginal cost is

$$c_{sbg} = \frac{w}{\phi_{s,g}} = \frac{w h(g)}{\phi_s}. \tag{60}$$

Under monopolistic competition with elasticity  $\sigma > 1$ , optimal pricing is a constant markup over marginal cost:

$$p_{sbg} = \tilde{\sigma} c_{sbg}, \quad \tilde{\sigma} \equiv \frac{\sigma}{\sigma - 1}. \tag{61}$$

Substituting (60) and (61) into the revenue equation gives

$$\begin{aligned}
r_{sbg} &= E_b \zeta_{sbg} \left( \frac{\tilde{\sigma} w h(g)}{\phi_s P_b} \right)^{1-\sigma} \\
&= E_b \left( \frac{P_b}{\tilde{\sigma} w} \right)^{\sigma-1} \phi_s^{\sigma-1} h(g)^{-(\sigma-1)} \zeta_{sbg}.
\end{aligned} \tag{62}$$

Defining

$$D_b \equiv E_b \left( \frac{P_b}{\tilde{\sigma} w} \right)^{\sigma-1}, \tag{63}$$

we obtain

$$r_{sbg} = D_b \phi_s^{\sigma-1} h(g)^{-(\sigma-1)} \zeta_{sbg}. \tag{64}$$

Since variable operating profits under CES pricing are a constant share of revenue,

$$\pi_{sbg} = \frac{1}{\sigma} r_{sbg} = \frac{D_b}{\sigma} \phi_s^{\sigma-1} h(g)^{-(\sigma-1)} \zeta_{sbg}. \tag{65}$$

These are the revenue and operating-profit equations used in the supplier's product-activation and link-formation decisions.

## H.4 Active sets and the partial-equilibrium interpretation

The active sets in the preceding derivation are endogenous from the supplier's perspective but are taken as given when buyer  $b$  allocates expenditure. The supplier first chooses firm-level product scope  $G_s$  and evaluates buyer relationships. Conditional on an active relationship, product–buyer match realizations determine the set  $\mathcal{G}_{sb} \subseteq \{1, \dots, G_s\}$  of products that are activated for buyer  $b$ . Products outside this set have  $q_{sbg} = r_{sbg} = 0$  and do not enter the buyer's CES price index.

The same match shifter  $\zeta_{sbg}$  therefore affects both the profitability of activating a product for a particular buyer and, conditional on activation, the product's realized expenditure share. This is why the benchmark matching model uses the same source of product–buyer heterogeneity for selection into trade and for sales conditional on trade.

From the perspective of an individual supplier, the buyer-side demand environment summarized by  $D_b$  is taken as given. In the underlying atomless-supplier environment, the realization of supplier  $s$ 's own product–buyer match  $\zeta_{sbg}$  has a negligible effect on buyer  $b$ 's price index  $P_b$ , and therefore does not feed back into  $D_b$ . The supplier can consequently evaluate expected product-level profitability conditional on  $(\phi_s, D_b)$  while holding the buyer demand shifter fixed.

With the normalization  $\mathbb{E}[\zeta_{sbg}] = 1$ , expected operating profits prior to conditioning on product

activation are therefore

$$\mathbb{E}[\pi_{sbj}(\phi_s, \zeta_{sbj}) \mid \phi_s, D_b] = \frac{D_b}{\sigma} \phi_s^{\sigma-1} h(g)^{-(\sigma-1)}. \quad (66)$$

This expectation is used when evaluating the profitability of products and buyer relationships in the supplier's problem.

The theoretical demand system satisfies buyer-level expenditure exhaustion by construction. In the quantitative implementation, however, the buyer-side profitability shifter  $D_b$  is treated as a reduced-form primitive and is not recomputed from the realized finite set of simulated suppliers. The finite supplier population is therefore interpreted as an approximation to an underlying atomless supplier environment. Accordingly, the quantitative simulation need not impose the buyer-level accounting identity

$$\sum_{s \in \mathcal{S}_b} \sum_{g \in \mathcal{G}_{sb}} r_{sbj} = E_b \quad (67)$$

exactly for every simulated buyer. An individual supplier instead takes  $D_b$  and the buyer-side demand environment as given when solving its product-scope, product-activation, and link-formation problems.

## H.5 Regularity conditions

Supplier productivity follows a Pareto distribution,

$$\phi_s \sim \text{Pareto}(\theta_\phi, \phi_{\min}), \quad \phi_s \geq \phi_{\min} > 0. \quad (68)$$

For a Pareto random variable,

$$\mathbb{E}[\phi_s^k] < \infty \quad \Longleftrightarrow \quad \theta_\phi > k. \quad (69)$$

In the matching model, the product–buyer match shock is also Pareto distributed,

$$\zeta_{sbj} \sim \text{Pareto}(\theta_\zeta, \zeta_{\min}), \quad \theta_\zeta > 1, \quad (70)$$

with  $\zeta_{\min}$  chosen so that  $\mathbb{E}[\zeta_{sbj}] = 1$ .

As shown in the characterization of expected product surplus, the relevant interior-Pareto term scales with supplier productivity as

$$\phi_s^{\theta_\zeta(\sigma-1)}. \quad (71)$$

Existence of the corresponding population moment therefore requires

$$\theta_\phi > \theta_\zeta(\sigma - 1). \quad (72)$$

This condition is stronger than the standard requirement  $\theta_\phi > \sigma - 1$  that applies in the absence of the additional selection induced by product–buyer matching.

Condition (72) is sufficient but not necessary when developed scope is bounded. For any finite  $G_{\max}$ , sufficiently productive suppliers have every product in the region where activation is certain, where expected surplus is linear in  $\phi_s^{\sigma-1}$ . The relevant population moments then exist whenever  $\theta_\phi > \sigma - 1$ . In the calibration,  $\theta_\phi = 4.63$  is estimated through  $\theta_\phi = \theta_\zeta(\sigma - 1) + g$  with  $g > 0$ , so the stronger condition (72) also holds.

## H.6 Summary of model primitives

Table 12: Model primitives

| Class          | Object        | Role                                  |
|----------------|---------------|---------------------------------------|
| Supplier       | $\phi_s$      | Supplier productivity                 |
| Supplier       | $c_s$         | Supplier networking cost              |
| Buyer          | $D_b$         | Buyer profitability shifter           |
| Buyer          | $\bar{f}_b^L$ | Average link-formation difficulty     |
| Buyer          | $\bar{f}_b^A$ | Average product-activation difficulty |
| Supplier–buyer | $\eta_{sb}$   | Pair-specific cost heterogeneity      |
| Product–buyer  | $\zeta_{sbq}$ | Match quality and realized appeal     |

*Notes:* The matching model uses  $\zeta_{sbq}$  both to determine product–buyer selection and to generate realized sales dispersion conditional on trade. The deterministic common-ladder benchmark, as in [Arkolakis et al. \(2021\)](#), instead imposes a common product ranking across buyers and therefore has no product–buyer match shock. We allow a separate ex-post appeal shock  $\xi_{sbq}$ , normalized so that  $\mathbb{E}[\xi_{sbq}] = 1$ , to generate residual product-level sales dispersion in this benchmark.

The primitive sources of heterogeneity satisfy

$$\phi_s \perp D_b \perp \eta_{sb} \perp c_s \perp \zeta_{sbq}, \quad (73)$$

subject to the buyer-level heterogeneity encoded in  $(D_b, \bar{f}_b^L, \bar{f}_b^A)$ .

Although the underlying shocks in (73) are mutually independent, the realized link-formation and product-activation costs are not independent because both contain the common components  $\eta_{sb}$  and  $c_s$ . This shared structure allows supplier–buyer pairs that are easier to establish to also be easier to operate along the product margin.

# I Product–Buyer Matching and Expected Surplus

Fix supplier scope  $G_s$  and define

$$A_{sb} \equiv \frac{D_b}{\sigma} \phi_s^{\sigma-1}.$$

Product  $g$  is activated for buyer  $b$  whenever

$$\zeta_{sb} \geq z_{sb}(g) \equiv \frac{f_{sb,0}^A}{A_{sb}} g^\kappa, \quad \kappa \equiv \mu + \alpha(\sigma - 1).$$

Given the Pareto distribution of  $\zeta$ , the activation probability is

$$p_{sb}(g) \equiv \Pr(\zeta \geq z_{sb}(g)) = \min \left\{ 1, \left( \frac{\zeta_{\min}}{z_{sb}(g)} \right)^{\theta_\zeta} \right\}. \quad (74)$$

Expected net surplus from product  $g$  is

$$\Psi_{sb}(g) = \mathbb{E} \left[ \left( A_{sb} g^{-\alpha(\sigma-1)} \zeta - f_{sb,0}^A g^\mu \right)^+ \right]. \quad (75)$$

Evaluating the Pareto integral gives

$$\Psi_{sb}(g) = \begin{cases} A_{sb} g^{-\alpha(\sigma-1)} - f_{sb,0}^A g^\mu, & z_{sb}(g) \leq \zeta_{\min}, \\ \frac{f_{sb,0}^A g^\mu p_{sb}(g)}{\theta_\zeta - 1}, & z_{sb}(g) > \zeta_{\min}. \end{cases} \quad (76)$$

For products in the interior Pareto region,  $z_{sb}(g) > \zeta_{\min}$ , this becomes

$$\Psi_{sb}(g) = \psi_{sb} g^{\mu - \kappa \theta_\zeta}, \quad (77)$$

where

$$\psi_{sb} = \frac{\zeta_{\min}^{\theta_\zeta}}{\theta_\zeta - 1} A_{sb}^{\theta_\zeta} (f_{sb,0}^A)^{1-\theta_\zeta}. \quad (78)$$

# J Numerical Solution

The quantitative model is solved using the exact discrete supplier problem. For each supplier  $s$  and candidate scope  $G \in \{0, 1, \dots, G_{\max}\}$ , we first compute

$$C_{sb}(G) = \sum_{g=1}^G \Psi_{sb}(g)$$

for every potential buyer. Expected net value from buyer  $b$  is

$$V_{sb}(G) = [C_{sb}(G) - f_{sb}^L]^+.$$

Supplier profits at candidate scope  $G$  are therefore

$$\Pi_s(G) = \sum_{b \in \mathcal{B}} V_{sb}(G) - F^P(G). \quad (79)$$

Optimal product scope is

$$G_s^* \in \arg \max_G \Pi_s(G). \quad (80)$$

Given  $G_s^*$ , the link-formation rule determines  $\mathcal{B}_s^*$ . Conditional on an active link, product–buyer match shocks are drawn and determine the realized set  $\mathcal{G}_{sb}^*$  and corresponding sales.

Existence follows from the finite product-scope choice set. The closed-form expression in Section 4.5 is used for economic interpretation; all quantitative results use the exact discrete solution.

## J.1 Distributions and simulation design

The quantitative model draws all primitives from parametric distributions. Supplier productivity is Pareto,

$$\phi_s = \phi_{\min} (1 - u_s)^{-1/\theta_\phi}, \quad u_s \in \left\{ \frac{i-1/2}{S} : i = 1, \dots, S \right\}, \quad \theta_\phi = \theta_\zeta(\sigma - 1) + g, \quad g > 0,$$

so that the productivity tail always satisfies the condition in equation (72); the gap  $g$  is estimated.

Supplier networking cost shifters, buyer demand shifters, and buyer-level cost scales are log-normal, each normalized to have the stated mean:

$$\ln c_s \sim \mathcal{N}\left(-\frac{\omega_c^2}{2}, \omega_c^2\right), \quad \ln D_b = \omega z_b - \frac{\omega^2}{2}, \quad z_b = \Phi^{-1}\left(\frac{b-1/2}{B}\right),$$

$$\bar{f}_b^L = \bar{f}^L \epsilon_b^L, \quad \bar{f}_b^A = \bar{f}_0^A \epsilon_b^A, \quad \ln \epsilon_b^j \sim \mathcal{N}\left(-\frac{\omega_{fj}^2}{2}, \omega_{fj}^2\right), \quad j \in \{L, A\},$$

so that  $\mathbb{E}[c_s] = \mathbb{E}[D_b] = \mathbb{E}[\epsilon_b^L] = \mathbb{E}[\epsilon_b^A] = 1$  and  $\bar{f}^L$  and  $\bar{f}_0^A$  are the mean link and activation cost scales. The pair-specific cost shock is lognormal,

$$\ln \eta_{sb} \sim \mathcal{N}\left(-\frac{\omega_\eta^2}{2}, \omega_\eta^2\right),$$

drawn independently for every supplier–buyer pair. The same draw enters both costs of the pair,

$f_{sb}^L = \bar{f}_b^L \eta_{sb} c_s$  and  $f_{sb,0}^A = \bar{f}_b^A \eta_{sb} c_s$ . Product–buyer match qualities are Pareto,

$$\zeta_{sbg} = \zeta_{\min} (1 - v_{sbg})^{-1/\theta_\zeta}, \quad v_{sbg} \sim U(0, 1), \quad \zeta_{\min} = \frac{\theta_\zeta - 1}{\theta_\zeta},$$

drawn independently for each active link and each developed product after the supplier’s scope and link decisions. They therefore affect realized selection and sales but not the ex-ante decisions, which depend only on the expected surplus in equation (22). To reduce simulation noise, supplier productivities and buyer qualities are placed on deterministic quantile grids of their distributions (Pareto and lognormal, respectively) and randomly assigned to firms; all other shocks are random draws, independent of each other and of the grids.

**Buyer pool and simulation.** The simulated economy consists of  $S$  suppliers and a common pool of  $B$  potential buyers. Buyer characteristics  $(D_b, \bar{f}_b^L, \bar{f}_b^A)$  are drawn once and are the same for all suppliers, so buyers differ in how attractive they are to every supplier, while  $\eta_{sb}$  and  $\zeta_{sbg}$  generate pair- and product-specific heterogeneity. Realized match qualities are then drawn on active links, and all moments are computed from realized trade. We use  $S = 8,000$  suppliers,  $B = 16,000$  potential buyers and  $G_{\max} = 300$ . The parameters are estimated on two independent simulated economies with fixed random seeds, so the SMM objective is smooth and reproducible; all reported results evaluate the calibrated parameters on four economies, the two used in estimation and two additional independent draws.

## K Additional diagnostics for buyer–product co-movement

This appendix provides additional diagnostics for the relationship between buyer base and product scope in the calibrated matching model. The main text shows that product scope and buyer base remain positively related after conditioning on supplier productivity. Here we report two complementary exercises.

First, Table 13 decomposes the variation in log product scope into the portion associated with supplier productivity and the additional portion associated with buyer base. Supplier productivity alone explains 59 percent of the cross-sectional variation in log product scope. Conditional on productivity, buyer base accounts for most of the remaining variation: adding  $\log N_s$  raises the explained share from 0.588 to 0.900, about 76 percent of the residual variation. Networking-cost heterogeneity accounts for a smaller part: adding  $\log c_s$  instead raises the explained share to 0.864.

Second, Table 14 reports the correlation between product scope and buyer base separately within productivity quintiles. The relationship remains strongly positive within every quintile, both in levels and in logs. This confirms that the positive co-movement documented in the main text is not driven solely by sorting of more productive suppliers into both larger product portfolios

Table 13: Variance decomposition of  $\log \widehat{G}_s$  (calibrated model)

|                                                                             | Share of $\text{Var}(\log \widehat{G}_s)$ |
|-----------------------------------------------------------------------------|-------------------------------------------|
| Explained by $\log \phi_s$ alone                                            | 0.588                                     |
| Explained by $\log \phi_s + \log N_s$                                       | 0.900                                     |
| Explained by $\log \phi_s + \log c_s$                                       | 0.864                                     |
| Within- $\phi$ residual share                                               | 0.412                                     |
| Additional share explained by $\log N_s$<br>as % of within- $\phi$ residual | 0.312<br>75.7%                            |

*Notes:* Multiproduct suppliers, four simulated economies at the calibrated parameters. Productivity enters as a cubic in  $\log \phi_s$ ;  $c_s$  is the supplier's networking cost.

and larger buyer bases.

Table 14: Within- $\phi$ -quintile correlations between  $\widehat{G}_s$  and  $N_s$  (calibrated model)

| Quintile | $\rho(\widehat{G}_s, N_s)$ | $\rho(\log \widehat{G}_s, \log N_s)$ |
|----------|----------------------------|--------------------------------------|
| Q1       | 0.837                      | 0.817                                |
| Q2       | 0.861                      | 0.857                                |
| Q3       | 0.880                      | 0.884                                |
| Q4       | 0.898                      | 0.914                                |
| Q5       | 0.820                      | 0.967                                |
| Pooled   | 0.763                      | 0.947                                |

*Notes:* Multiproduct suppliers, pooled across four simulated economies at the calibrated parameters.

## K.1 How broadly is the return to product adoption spread across buyers?

The match-adjusted buyer-base aggregator in equation (34) places heterogeneous weights on active buyers. Let

$$w_{sb} \equiv \frac{D_b^{\theta_\zeta} (f_{sb,0}^A)^{1-\theta_\zeta}}{\Lambda_s^M}, \quad b \in \mathcal{B}_s,$$

denote buyer  $b$ 's share of the aggregator, so that  $\sum_{b \in \mathcal{B}_s} w_{sb} = 1$ . To summarize how concentrated these contributions are, we use the effective number of buyers, the inverse Herfindahl index of these weights:

$$N_s^{\text{eff}} \equiv \left( \sum_{b \in \mathcal{B}_s} w_{sb}^2 \right)^{-1}. \quad (81)$$

The effective number equals  $N_s$  when all active buyers contribute equally and approaches one when a single buyer dominates.

Table 15 shows that the return to product adoption is spread across many buyer relationships. The effective number of buyers rises steadily with the active buyer count, from about 6 among suppliers with 8 to 15 buyers to about 79 among suppliers with more than 256, and the largest buyer accounts for only 6 percent of the aggregator in that group. The median supplier has about 8 effective buyers. The small amortization share in Section 7.1 therefore does not arise because the return to a new product is concentrated in a few buyers. It arises because of the structure of adoption costs: developing a product is cheap at the margin but the cost rises steeply with scope, so the developed portfolio is pinned down mainly by productivity and networking costs, which explain 94 percent of its variation, rather than by the number of buyers. Additional buyers do raise the return to marginal products, but those products are peripheral and rarely sold, so the response barely reaches observed scope.

Table 15: Effective number of buyers in the product-adoption aggregator


| Active buyers $N_s$ | Median $N_s$ | Median $N_s^{\text{eff}}$ | $N_s^{\text{eff}}/N_s$ | Top-1 share |
|---------------------|--------------|---------------------------|------------------------|-------------|
| 1                   | 1            | 1.00                      | 1.000                  | 1.000       |
| 2–3                 | 3            | 1.92                      | 0.781                  | 0.647       |
| 4–7                 | 5            | 3.43                      | 0.645                  | 0.432       |
| 8–15                | 11           | 5.96                      | 0.547                  | 0.297       |
| 16–31               | 22           | 10.21                     | 0.471                  | 0.208       |
| 32–63               | 42           | 16.90                     | 0.396                  | 0.152       |
| 64–127              | 86           | 27.94                     | 0.323                  | 0.112       |
| 128–255             | 168          | 43.39                     | 0.254                  | 0.088       |
| 256+                | 460          | 78.89                     | 0.165                  | 0.063       |
| All firms           | 16           | 8.14                      | 0.487                  | 0.242       |
| Link-weighted       | 855          | 70.50                     | 0.251                  | 0.113       |

*Notes:* Pooled across four simulated economies at the calibrated parameters. Active (ex-ante) links. Weights  $w_{sb} \propto D_b^{\theta\zeta} (f_{sb,0}^A)^{1-\theta\zeta}$ , normalized to sum to one within supplier;  $N_s^{\text{eff}} = (\sum_b w_{sb}^2)^{-1}$ . Columns report medians within each group. The link-weighted row weights suppliers by  $N_s$ .

## K.2 Sales across margins

Figure 25 relates simulated supplier sales to buyer base, product scope, and supplier productivity. Sales increase with all three objects, but the relationships differ in tightness. Buyer count is the most tightly associated with sales (correlation 0.91), followed by product scope (0.88) and productivity (0.84). This ordering is consistent with Table 8: firm size in the model is primarily a matter of how many buyers a supplier reaches. At a given buyer count or product scope, more productive suppliers earn more, so the productivity quintiles fan out; at a given productivity, sales still span several log points because suppliers differ in networking costs and in the buyers they reach. These patterns mirror the empirical fact that larger suppliers operate with broader buyer and product scopes (Figure 1).

Figure 25: Total sales across margins (calibrated model)



Notes: Log total sales  $\log T_s$  against log buyer count  $\log N_s$  (panel a), log product scope  $\log \hat{G}_s$  (panel b), and log productivity  $\log \phi_s$  (panel c), colored by productivity quintile. One simulated economy, the same as in Figure 8: 7,696 suppliers with at least one product and one buyer. Correlations with  $\log T_s$ : 0.91, 0.88 and 0.84.

## L Alternative calibrations

This appendix reports two alternative specifications: a calibration with a less convex adoption cost and a common-product-ladder benchmark. All are evaluated on two simulated economies, so moments can differ slightly from the main tables, which average four.

**Fixing the adoption-cost convexity.** The curvature  $\delta$  governs how easily the developed portfolio can expand in response to a larger buyer base, and therefore how much room the model leaves for amortization. Table 16 recalibrates the model with  $\delta$  fixed at 1.5 instead of the estimated 2.74. A less convex adoption cost lets the developed portfolio respond more to buyers (29 instead of 21 percent of  $\beta$ ), but amortization still accounts for less than 5 percent of the co-movement in observed scope, and the fit and the fixed-cost comparative statics are nearly unchanged.

Table 16: Robustness: fixing the adoption-cost convexity

|                                     | $\delta$ estimated (2.74) | $\delta = 1.5$  |
|-------------------------------------|---------------------------|-----------------|
| $\beta$ (data 0.555)                | 0.518                     | 0.540           |
| $\beta$ , developed portfolio fixed | 0.513                     | 0.515           |
| Amortization share of $\beta$       | 1.0% (2.2)                | 4.5% (2.2)      |
| Developed-portfolio response        | 21%                       | 29%             |
| $F^P$ -50%: products / buyers       | +7.0% / +1.8%             | +4.9% / +1.8%   |
| $\bar{f}^L$ -50%: products / buyers | +14.4% / +45.0%           | +14.0% / +44.6% |
| SMM criterion (17 targets)          | 0.47                      | 0.48            |

**Common product ladder.** Appendix N calibrates the deterministic-ladder benchmark to the same data. Table 17 compares it with the matching model.

Table 17: Benchmark: a common product ladder

|                                | Data  | Common ladder | Matching |
|--------------------------------|-------|---------------|----------|
| Criterion, 14 shared targets   |       | 1.45          | 0.36     |
| Scope-buyer elasticity $\beta$ | 0.555 | 0.13          | 0.51     |
| $\beta$ with $F^P = 0$         |       | 0.11          | 0.59     |
| Coverage (mean)                | 0.61  | 1.00          | 0.58     |
| Non-nested buyer pairs         | 0.27  | 0.00          | 0.57     |
| Top product missing from link  | 0.36  | 0.03          | 0.24     |

## M Extension: random product appeal

In the benchmark model, a supplier’s products differ systematically through their rank  $g$  and idiosyncratically through the buyer-specific match  $\zeta_{sb g}$ . Products are developed in rank order and each new product is less appealing than the previous one, so the marginal developed product is peripheral and rarely sold. This limits how much a larger buyer base can raise observed scope through product development, and makes the benchmark a lower bound on amortization.

This appendix relaxes that structure in two ways. First, each product carries its own appeal  $\varepsilon_{sg}$ , known to the supplier and common to all of its buyers, so that product revenue becomes

$$r_{sb g} = \tilde{D}_b \phi_s^{\sigma-1} g^{-\alpha(\sigma-1)} \varepsilon_{sg} \zeta_{sb g}, \quad \ln \varepsilon_{sg} \sim \mathcal{N}\left(-\frac{\sigma_\varepsilon^2}{2}, \sigma_\varepsilon^2\right). \quad (82)$$

Products are still developed in the order  $g = 1, 2, \dots$  at rising adoption cost, but a newly developed product can now be more attractive than earlier ones, and its appeal is correlated across all of the supplier’s buyers. Second, we allow a buyer’s demand shifter to enter as  $\tilde{D}_b \propto D_b^\gamma$ , normalized to mean one, with  $\gamma$  estimated;  $\gamma = 1$  in the benchmark. When  $\gamma < 1$ , a buyer’s contribution to the return to a new product depends less on its size, so the firm amortizes over the number of buyers rather than over their total size.

We report three calibrations. The first adds random appeal and  $\gamma$  to the benchmark targets. The other two also target the static analog of the incumbent-buyer coefficient of Appendix D, with  $\sigma_\varepsilon$  estimated or fixed at 1.5. The model counterpart of that coefficient compares two equilibria of the same economy, before and after halving every link cost with match draws held fixed, and divides the change in log products sold to the buyers served before the shock by the change in log buyers. Because the model is static, it cannot contain the dynamics behind the data regression, such as relationships that deepen over time or adoption lags, and no calibration comes close to the empirical value of 0.38.

Table 18 reports the results. The targeted moments are consistent with product appeal: estimating  $\sigma_\varepsilon$  (0.70) and  $\gamma$  improves the fit on the 17 targets from 0.413 to 0.223, with two additional parameters. Calibrations that fit the data comparably well imply an amortization share between

6 and 12 percent of  $\beta$ , rising with the dispersion of appeal from 6.3 percent at  $\sigma_\varepsilon = 0.70$  to 11.7 percent at  $\sigma_\varepsilon = 1.15$ . The targeted moments do not determine  $\sigma_\varepsilon$  precisely within this range. A calibration with more dispersed appeal still ( $\sigma_\varepsilon = 1.5$ ) implies 15.9 percent, but its scope-buyer elasticity is 0.44 against 0.555 in the data, 22 percent below, so we treat it as an upper bound rather than as part of the range. We keep the benchmark without product appeal as our main specification because it is more parsimonious and all other quantitative results in the paper are computed from it.

Two caveats apply. First, the shares are computed with  $\beta$  controlling for the supplier’s average product appeal, the analog of a supplier fixed effect when appeal is random; with the controls of equation (39) alone, which omit it, the amortization share is about 1 percent in every calibration. Second, with random appeal the counterfactual portfolio predicted from productivity and networking costs fits the actual one less well (an  $R^2$  of 0.73 to 0.81, against 0.94 in the benchmark) and is somewhat smaller (mean log scope falls by 0.04 to 0.06), so part of the fall in  $\beta$  may reflect smaller portfolios rather than amortization. Even under these assumptions, matching accounts for the large majority of the co-movement. None of the calibrations comes close to the empirical incumbent-buyer coefficient (at most 0.06 against 0.38), which points again to dynamic forces outside the static model.

Table 18: Random product appeal and amortization

|                                                | Benchmark  | Random appeal<br>( $\sigma_\varepsilon$ estimated) | More dispersed appeal<br>(incumbent targeted) | Upper bound<br>( $\sigma_\varepsilon = 1.5$ ) |
|------------------------------------------------|------------|----------------------------------------------------|-----------------------------------------------|-----------------------------------------------|
| Product-appeal dispersion $\sigma_\varepsilon$ | 0          | 0.703                                              | 1.147                                         | 1.500                                         |
| Buyer-size elasticity $\gamma$                 | 1          | 0.805                                              | 0.788                                         | 0.860                                         |
| $\beta$ (data 0.555)                           | 0.518      | 0.558                                              | 0.513                                         | 0.442                                         |
| Incumbent coefficient (data 0.38)              | 0.030      | 0.044                                              | 0.064                                         | 0.060                                         |
| Products per link (data 3.24)                  | 3.47       | 3.53                                               | 4.42                                          | 3.71                                          |
| Fit, 17 targets                                | 0.413      | 0.223                                              | 0.299                                         | 0.208                                         |
| Amortization share of $\beta$                  | 1.0% (2.2) | 6.3% (1.3)                                         | 11.7% (1.2)                                   | 15.9% (1.2)                                   |

*Notes:* Calibrated on two simulated economies and evaluated on four; standard errors across economies in parentheses. The fit is computed at the calibration stage, on two economies, so the benchmark value differs from the four-economy value in Table 4. The last two columns also target the incumbent-buyer coefficient (weight 2), a static proxy computed from two equilibria before and after halving link costs. Amortization share =  $1 - \beta_{\text{fixed}}/\beta$  from the fixed-portfolio counterfactual of Section 7.1; in the random-appeal columns,  $\beta$  also controls for the supplier’s average product appeal.

## N Deterministic-Ladder Benchmark (M1)

This appendix presents the deterministic-ladder benchmark, M1, used in the quantitative comparisons in Section 5. M1 preserves the same basic supplier productivity, buyer heterogeneity, fixed-cost structure, and firm-level product-adoption decision as the matching model in the main

text. The key difference is the mapping from the supplier's developed portfolio into products sold to individual buyers.

In M1, products follow a common deterministic profitability ranking across buyers. Conditional on serving buyer  $b$ , supplier  $s$  sells a prefix of its firm-level product ladder,

$$\{1, \dots, G_{sb}\}, \quad G_{sb} \leq G_s.$$

Thus, buyers can differ in how far they extend down the supplier's product ranking, but they cannot purchase non-nested subsets of products. There is no product–buyer match shock  $\zeta_{sb}$ . Instead, realized product-level sales contain an ex-post appeal shock  $\xi_{sb}$  with  $\mathbb{E}[\xi_{sb}] = 1$ . Because  $\xi_{sb}$  is realized after the scope and link decisions, it generates residual sales dispersion but does not affect which products are activated or which links are formed.

## N.1 Stage 2: Deterministic link-level product scope

Fix supplier  $s$ 's firm-level product scope  $G_s$ . As in the baseline model, define the buyer–supplier profitability shifter

$$A_{sb}(\phi_s) \equiv \frac{D_b}{\sigma} \phi_s^{\sigma-1}. \quad (83)$$

Products are ranked according to

$$h(g) = g^\alpha,$$

and the buyer-specific activation cost is

$$f_{sb}^A(g) = f_{sb,0}^A g^\mu, \quad f_{sb,0}^A = \bar{f}_b^A \eta_{sb} c_s.$$

Because the ex-post appeal shock has mean one and does not affect decisions, the expected operating profit generated by rank  $g$  before paying its buyer-specific activation cost is

$$A_{sb}(\phi_s) g^{-\alpha(\sigma-1)}. \quad (84)$$

The incremental gross payoff from activating rank  $g$  on link  $(s, b)$  is therefore

$$\Delta u_{sb}(g; \phi_s) = A_{sb}(\phi_s) g^{-\alpha(\sigma-1)} - f_{sb,0}^A g^\mu. \quad (85)$$

Define, as in the main text,

$$\kappa \equiv \mu + \alpha(\sigma - 1). \quad (86)$$

When  $\kappa > 0$ , the ratio of activation cost to operating profit rises with rank, so the incremental payoff in (85) crosses zero at most once. The unconstrained buyer-specific product cutoff is

therefore

$$\begin{aligned}
\tilde{G}_{sb}(\phi_s) &= \max \{g : A_{sb}(\phi_s)g^{-\alpha(\sigma-1)} \geq f_{sb,0}^A g^\mu\} \\
&= \left\lfloor \left( \frac{A_{sb}(\phi_s)}{f_{sb,0}^A} \right)^{1/\kappa} \right\rfloor \\
&= \left\lfloor \left( \frac{D_b \phi_s^{\sigma-1}}{\sigma \bar{f}_b^A \eta_{sb} c_s} \right)^{1/\kappa} \right\rfloor.
\end{aligned} \tag{87}$$

Conditional on the firm's developed scope  $G_s$ , the desired product depth of buyer  $b$  is

$$G_{sb}^0(G_s; \phi_s) = \min \{G_s, \tilde{G}_{sb}(\phi_s)\}. \tag{88}$$

Hence every active buyer receives a deterministic prefix of the supplier's product hierarchy.

Before paying the link fixed cost, the gross value of the relationship is

$$U_{sb}^{\text{M1}}(G_s; \phi_s) = \sum_{g=1}^{G_{sb}^0(G_s; \phi_s)} [A_{sb}(\phi_s)g^{-\alpha(\sigma-1)} - f_{sb,0}^A g^\mu]. \tag{89}$$

Define net link surplus as

$$S_{sb}^{\text{M1}}(G_s; \phi_s) \equiv U_{sb}^{\text{M1}}(G_s; \phi_s) - f_{sb}^L. \tag{90}$$

The supplier activates the relationship whenever

$$I_{sb}^{\text{M1}}(G_s; \phi_s) = \mathbf{1} \{S_{sb}^{\text{M1}}(G_s; \phi_s) \geq 0\}. \tag{91}$$

Thus the realized link-level scope is

$$G_{sb} = I_{sb}^{\text{M1}}(G_s; \phi_s) G_{sb}^0(G_s; \phi_s). \tag{92}$$

Conditional on trade, realized product-level sales are

$$r_{sb}^{\text{M1}} = D_b \phi_s^{\sigma-1} g^{-\alpha(\sigma-1)} \xi_{sb}, \quad g \leq G_{sb}, \tag{93}$$

where  $\mathbb{E}[\xi_{sb}] = 1$ . Unlike  $\zeta_{sb}$  in M2,  $\xi_{sb}$  affects realized sales only and plays no role in selection into product-buyer trade.

## N.2 Stage 1: Firm-level product scope

Anticipating the Stage 2 policy, supplier  $s$  chooses its firm-level scope  $G_s$  to maximize

$$\Pi_s^{\text{M1}}(G_s; \phi_s) = \sum_{b \in \mathcal{B}} [S_{sb}^{\text{M1}}(G_s; \phi_s)]^+ - F^{\text{P}}(G_s), \quad (94)$$

where

$$F^{\text{P}}(G_s) = \sum_{g=1}^{G_s} f^{\text{P}}(g), \quad f^{\text{P}}(g) = f_0^{\text{P}} g^\delta.$$

The exact optimal scope is therefore

$$G_s^{\star, \text{M1}} \in \arg \max_{G_s \in \{0, 1, \dots, G_{\max}\}} \Pi_s^{\text{M1}}(G_s; \phi_s). \quad (95)$$

The marginal value of increasing firm scope from  $g-1$  to  $g$  is

$$\Delta \Pi_s^{\text{M1}}(g) = \sum_{b \in \mathcal{B}} \left( [S_{sb}^{\text{M1}}(g)]^+ - [S_{sb}^{\text{M1}}(g-1)]^+ \right) - f^{\text{P}}(g). \quad (96)$$

This marginal value can again be separated into intensive and extensive components. Define

$$\mathcal{B}_s^{\text{bind}}(g) \equiv \left\{ b \in \mathcal{B}_s(g-1) : \tilde{G}_{sb} \geq g \right\}, \quad (97)$$

the previously active buyers for which rank  $g$  remains below the buyer-specific cutoff, and

$$\mathcal{B}_s^{\text{new}}(g) \equiv \mathcal{B}_s(g) \setminus \mathcal{B}_s(g-1), \quad (98)$$

the relationships activated when the supplier develops product  $g$ . Then

$$\begin{aligned} \Delta \Pi_s^{\text{M1}}(g) = & \underbrace{\sum_{b \in \mathcal{B}_s^{\text{bind}}(g)} [A_{sb} g^{-\alpha(\sigma-1)} - f_{sb}^A(g)]}_{\text{intensive: deepen existing prefixes}} \\ & + \underbrace{\sum_{b \in \mathcal{B}_s^{\text{new}}(g)} S_{sb}^{\text{M1}}(g)}_{\text{extensive: newly activated links}} - f^{\text{P}}(g). \end{aligned} \quad (99)$$

The important difference from M2 is the first term. In M1, buyer  $b$  contributes to the marginal return from product  $g$  only if  $\tilde{G}_{sb} \geq g$ . Thus, as the supplier moves farther down the product ladder, buyers progressively leave the amortization base.

### N.3 A scope ceiling in the deterministic ladder

The buyer-specific cutoffs imply an upper bound on firm-level product scope that is absent from the matching model.

[Scope ceiling in the deterministic-ladder model] Suppose  $f^P(g) > 0$  for products above the current portfolio. Then optimal firm-level scope in M1 satisfies

$$G_s^{*,M1} \leq \max_{b \in \mathcal{B}} \tilde{G}_{sb} = \max_{b \in \mathcal{B}} \left[ \left( \frac{D_b \phi_s^{\sigma-1}}{\sigma \bar{f}_b^A \eta_{sb} c_s} \right)^{1/\kappa} \right]. \quad (100)$$

*Proof.* Consider any  $G_s > \max_b \tilde{G}_{sb}$ . For every buyer, equation (88) implies

$$G_{sb}^0(G_s) = \tilde{G}_{sb}.$$

Increasing  $G_s$  further therefore changes neither the product set nor the gross value of any potential relationship. The first term of (94) is unchanged, while the firm continues to pay additional firm-level adoption costs. Hence such a value of  $G_s$  cannot be strictly optimal.  $\square$

This result clarifies an important limitation of the deterministic ladder. Amortization can raise product scope only while the marginal product is below the cutoff of at least one buyer. Once  $G_s$  reaches the largest buyer-specific cutoff, additional buyers' payoffs cannot be summed to support a product beyond that ceiling. The ceiling is therefore determined by the buyer with the deepest feasible product range rather than by the sum of surplus contributions across buyers.

At the same time, this feature generates a mechanical relationship between buyer breadth and product scope. A supplier with more buyers draws more buyer-specific cutoffs, increasing the probability that at least one buyer has a high  $\tilde{G}_{sb}$ . Firm scope can therefore increase with buyer count through the order statistic

$$\max_{b \in \mathcal{B}_s} \tilde{G}_{sb},$$

even in the absence of a firm-level cost to amortize. This is why the buyer–product relationship in M1 survives when  $F^P$  is set to zero. The matching model has an analogous non-amortization channel: with buyer-specific matches, observed scope is the union of products selected by at least one buyer (equation (36)), which also grows with the number of buyers. The fixed-portfolio counterfactual of Section 7.1 measures the size of this channel directly.

## N.4 Amortization in M1 and the activation-cost correction

M1 nevertheless contains a genuine amortization force below the scope ceiling. Holding the active and binding buyer sets locally fixed, the marginal condition for product  $g$  can be written as

$$\frac{\phi_s^{\sigma-1}}{\sigma} g^{-\alpha(\sigma-1)} \sum_{b \in \mathcal{B}_s^{\text{bind}}(g)} D_b \geq f_0^P g^\delta + c_s g^\mu \sum_{b \in \mathcal{B}_s^{\text{bind}}(g)} \bar{f}_b^A \eta_{sb}. \quad (101)$$

The first term on the right-hand side is the firm-level adoption cost, which is paid once. The second is different: activation costs are paid separately for every buyer to which the marginal product is supplied. Consequently, a larger buyer base raises both the return to adopting the product and the total activation cost associated with deploying it.

To make this correction explicit, define the activation-cost index at the marginal rank by

$$\Xi_s \equiv g^{\mu-\delta} \sum_{b \in \mathcal{B}_s^{\text{bind}}(g)} \bar{f}_b^A \eta_{sb}. \quad (102)$$

Then the local marginal condition can be expressed as

$$\frac{\phi_s^{\sigma-1}}{\sigma} \Lambda_s^{\text{bind}} g^{-\alpha(\sigma-1)} \geq (f_0^P + c_s \Xi_s) g^\delta, \quad \Lambda_s^{\text{bind}} \equiv \sum_{b \in \mathcal{B}_s^{\text{bind}}(g)} D_b. \quad (103)$$

Under the local proportional-buyer approximation—holding the composition of the relevant buyer set fixed while varying its size—the elasticity of scope with respect to buyer base is therefore attenuated relative to the formula that ignores buyer-specific activation costs:

$$\varepsilon_{G,N}^{\text{M1}} \simeq \underbrace{\frac{1}{\delta + \alpha(\sigma - 1)}}_{\text{clean amortization term}} \underbrace{\frac{f_0^P}{f_0^P + c_s \Xi_s}}_{\text{activation-cost correction}}. \quad (104)$$

The correction factor is strictly below one whenever buyer-specific activation costs are positive. The reason is economic: an additional buyer helps amortize the firm-level cost  $f_0^P$ , but it does not amortize the activation costs that must be paid again on each buyer relationship. Consequently, the expression  $1/[\delta + \alpha(\sigma - 1)]$  by itself overstates the strength of the amortization channel in M1.

Equation (104) should be interpreted as a local comparative-static approximation, not as a global characterization of M1. The exact solution is given by (95) and is additionally constrained by the buyer-specific scope ceiling in Proposition N.3.

## N.5 Marginal buyers and monotonicity

For any fixed firm-level scope  $G_s$ , buyers must be ordered by their net link surplus rather than by gross link value. Let

$$S_{sb}^{\text{M1}}(G_s; \phi_s) = U_{sb}^{\text{M1}}(G_s; \phi_s) - f_{sb}^L,$$

as in equation (90), and order potential buyers from highest to lowest according to this object. Let  $b(n; G_s)$  denote the buyer in position  $n$ . The number of active buyers is then

$$N_s(G_s; \phi_s) = \max \{n : S_{s,b(n; G_s)}^{\text{M1}}(G_s; \phi_s) \geq 0\}. \quad (105)$$

When an interior marginal buyer exists,

$$\begin{aligned} U_{s,b(N_s; G_s)}^{\text{M1}}(G_s; \phi_s) &\geq f_{s,b(N_s; G_s)}^L, \\ U_{s,b(N_s+1; G_s)}^{\text{M1}}(G_s; \phi_s) &< f_{s,b(N_s+1; G_s)}^L. \end{aligned} \quad (106)$$

Hence the relevant ordering for link formation is by net surplus, not by  $U_{sb}^{\text{M1}}$  alone.

The model generates the “more products  $\Rightarrow$  more buyers” direction weakly. Because all ranks up to  $\tilde{G}_{sb}$  have nonnegative incremental operating surplus,  $U_{sb}^{\text{M1}}(G_s)$  is weakly increasing in  $G_s$ . It follows that

$$\mathcal{B}_s^{\text{M1}}(G_s) \subseteq \mathcal{B}_s^{\text{M1}}(G_s + 1). \quad (107)$$

The reverse statement requires more care. For a product rank  $g$ , an additional buyer increases the gross marginal return to firm-level adoption only when that buyer itself supports rank  $g$ , that is, when

$$\tilde{G}_{sb} \geq g \iff A_{sb} g^{-\alpha(\sigma-1)} - f_{sb}^A(g) \geq 0. \quad (108)$$

Under this gross-payoff condition, adding the buyer weakly raises the left-hand side of the marginal adoption condition and therefore weakly raises the incentive to adopt product  $g$ . The condition cannot instead be imposed on arbitrary net link payoffs, because a potential buyer can have negative net surplus after the relationship fixed cost  $f_{sb}^L$  is paid and therefore never enter the amortization base.

This qualification is specific to the deterministic ladder. In M2, every active buyer contributes strictly positive expected surplus  $\Psi_{sb}(g) > 0$  for every finite product threshold, so the corresponding buyer contribution to the amortization sum does not disappear at a buyer-specific product ceiling.

## N.6 Comparison of the scope elasticities in M1 and M2

Ignoring the scope ceiling and buyer-specific activation-cost correction, the clean power-law approximation in M1 has scope exponent

$$\frac{1}{\delta + \alpha(\sigma - 1)}. \quad (109)$$

By contrast, the interior-Pareto approximation in M2 has exponent

$$\frac{1}{\delta + \kappa\theta_\zeta - \mu}. \quad (110)$$

Using  $\kappa = \mu + \alpha(\sigma - 1)$ , the denominator in M2 can be written as

$$\delta + \kappa\theta_\zeta - \mu = \delta + \alpha(\sigma - 1) + \kappa(\theta_\zeta - 1). \quad (111)$$

Since  $\theta_\zeta > 1$  and  $\kappa > 0$ , the exponent on the buyer-base aggregator in the M2 closed form is smaller than the corresponding clean M1 exponent when the remaining primitives are held fixed. This does not imply that amortization is weaker in M2. The two expressions are elasticities with respect to different objects. M1 depends on the raw buyer-profitability sum for buyers whose cutoff still reaches the marginal rank, while M2 depends on the match-adjusted expected-surplus aggregator. Moreover, the clean M1 expression abstracts from both the activation-cost correction in equation (104) and the scope ceiling in Proposition N.3. The two exponents therefore cannot be interpreted as a direct structural ranking of the strength of complementarity across the two models.

## N.7 Cost accounting and why M1 and M2 are calibrated separately

There is one additional accounting difference between the two specifications. In M1, if buyer  $b$  receives a prefix of length  $G_{sb}$ , the supplier pays activation costs on every product in that prefix:

$$F_{sb}^{A,M1}(G_{sb}) = \sum_{g=1}^{G_{sb}} f_{sb}^A(g). \quad (112)$$

In M2, by contrast, activation costs are paid only on products whose realized product–buyer match is sufficiently strong to make them active:

$$F_{sb}^{A,M2}(\mathcal{G}_{sb}) = \sum_{g \in \mathcal{G}_{sb}} f_{sb}^A(g). \quad (113)$$

At common finite parameter values, M2 therefore avoids activation costs for developed products that fail to match a particular buyer, while M1 pays the cost of every product in the deterministic prefix.

For this reason, parameter values should not be compared mechanically across M1 and M2. The two models are calibrated separately to their respective moment sets, and their quantitative implications are evaluated at their own calibrated parameter values. The purpose of M1 is not to provide a nested same-parameter counterfactual to M2, but to show what the common deterministic product ladder implies when it is given its own best quantitative fit.

Taken together, these differences explain why M1 is useful as a benchmark but not as the preferred model. The deterministic ladder preserves the economic force that firm-level product costs can be amortized across buyers, but it also imposes common prefix product sets, introduces a buyer-specific scope ceiling, and generates a mechanical deepest-buyer channel through which buyer breadth and product scope can co-move even when the firm-level adoption cost is removed. M2 replaces the deepest-buyer channel with buyer-specific matching, which the data require (coverage below one, non-nested product sets) and whose contribution relative to amortization can be measured.

## O Extension 1: Nested CES and within-link cannibalization

The baseline matching model in Section 4 abstracts from within-link cannibalization: conditional on realized product–buyer matches, adding a product on link  $(s, b)$  does not reduce the profitability of the supplier’s other products sold to the same buyer. This separability follows from the single CES demand system, under which a supplier’s own products and the products of other suppliers are equally substitutable.

In this section, we relax that assumption by introducing a nested CES demand system, with an outer nest across suppliers, governed by elasticity  $\nu$ , and an inner nest across products within each supplier, governed by elasticity  $\varepsilon$ . To isolate the cannibalization mechanism tractably, we write the within-link scope problem in the deterministic rank-based form of Appendix N: conditional on serving buyer  $b$ , supplier  $s$  chooses how far to extend down its product ranking on that relationship. The extension should therefore be interpreted as a robustness exercise for the separability assumption rather than as a full nested-CES version of the buyer-specific matching structure in Section 4. The firm-level adoption decision, buyer-link decision, and distinction between firm-level and link-level scope are preserved. The new ingredient is that products sold by the same supplier to the same buyer interact through the inner CES nest, generating within-link cannibalization whenever  $\varepsilon \neq \nu$ .

## O.1 Demand: nested CES across suppliers and products

Let  $\mathcal{S}_b \subseteq \mathcal{S}$  denote the endogenous set of suppliers active on buyer  $b$ , and let  $\mathcal{G}_{sb} \subseteq \{1, \dots, G_s\}$  denote the set of products that supplier  $s \in \mathcal{S}_b$  activates on link  $(s, b)$ , with cardinality  $G_{sb} = |\mathcal{G}_{sb}|$ . In the rank-based formulation used in this extension,  $\mathcal{G}_{sb} = \{1, \dots, G_{sb}\}$ . These objects are determined by the supplier's two-stage problem below and are taken as given when buyer  $b$  allocates expenditure. Products in the supplier's adopted portfolio that are not activated on the link, as well as suppliers that are not active on buyer  $b$ , do not enter the buyer's choice set.

**Outer nest: across suppliers.** Let  $X_{sb}$  denote the composite input that buyer  $b$  purchases from supplier  $s \in \mathcal{S}_b$ , the supplier's "variety" inside buyer  $b$ . The buyer's overall intermediate-input bundle is

$$M_b = \left[ \sum_{s \in \mathcal{S}_b} X_{sb}^{\frac{\nu-1}{\nu}} \right]^{\frac{\nu}{\nu-1}}, \quad \nu > 1, \quad (114)$$

where  $\nu$  governs substitution across suppliers: a higher  $\nu$  means that buyer  $b$  can more easily substitute expenditure across suppliers.

**Inner nest: across products within a supplier.** For each active supplier  $s \in \mathcal{S}_b$ , the within-supplier product-line composite is

$$X_{sb} = \left[ \sum_{g \in \mathcal{G}_{sb}} x_{sbg}^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}, \quad \varepsilon > 1, \quad (115)$$

where  $x_{sbg}$  is the quantity of product  $g$  that supplier  $s$  sells to buyer  $b$ , and  $\varepsilon$  governs substitution across products within a supplier's product line. When  $\varepsilon = \nu$ , the two nests collapse to the single-CES structure used in the separable benchmark, and the within-link cannibalization force described below disappears.

**Price indices.** Let  $p_{sbg}$  denote the delivered price of product  $g$  sold by supplier  $s$  to buyer  $b$ . The within-supplier, or inner-nest, price index is

$$P_{sb} = \left[ \sum_{g \in \mathcal{G}_{sb}} p_{sbg}^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}}, \quad (116)$$

which is dual to (115) and gives the minimum unit cost of one unit of the supplier composite  $X_{sb}$  given prices  $\{p_{sbg}\}_{g \in \mathcal{G}_{sb}}$ . The buyer-level, outer-nest price index across suppliers is

$$P_b = \left[ \sum_{s \in \mathcal{S}_b} P_{sb}^{1-\nu} \right]^{\frac{1}{1-\nu}}, \quad (117)$$

which is dual to (114). Appendix ?? derives the corresponding expenditure-minimization problem and revenue equations for the single-nest CES system used in the baseline model.

**Two-stage budgeting.** Given total expenditure  $E_b$ , buyer  $b$  first allocates spending across suppliers and then across products within each supplier:

$$R_{sb} \equiv P_{sb} X_{sb} = E_b \left( \frac{P_{sb}}{P_b} \right)^{1-\nu}, \quad (\text{outer}) \quad (118)$$

$$r_{sbg} \equiv p_{sbg} x_{sbg} = R_{sb} \left( \frac{p_{sbg}}{P_{sb}} \right)^{1-\varepsilon}, \quad g \in \mathcal{G}_{sb}, \quad (\text{inner}) \quad (119)$$

Equations (118)–(119) make the two effects of adding a product explicit. Adding a product to  $\mathcal{G}_{sb}$  lowers the inner-nest price index  $P_{sb}$ , which raises supplier  $s$ ’s share of buyer  $b$ ’s total expenditure through the outer nest. At the same time, expenditure within the supplier is reallocated across a larger set of products. This second force generates within-link cannibalization. Its importance is governed by the relative substitutability of products within the supplier and suppliers across the buyer’s input bundle. Cannibalization is strongest when products within a supplier are closer substitutes than different suppliers’ product lines, that is, when  $\varepsilon > \nu$ .

## O.2 Technology, pricing, and link-level revenue

The marginal-cost schedule is the same as in the benchmark:

$$c_{sbg} = \frac{w h(g)}{\phi_s}, \quad (120)$$

where  $w$  is the input price,  $\phi_s$  is supplier productivity, and  $h(g)$  is the rank schedule, increasing in  $g$ . Optimal pricing under monopolistic competition with outer-nest elasticity  $\nu$  gives a constant markup over marginal cost:<sup>8</sup>

$$p_{sbg} = \tilde{\nu} c_{sbg} = \tilde{\nu} \frac{w h(g)}{\phi_s}, \quad \tilde{\nu} \equiv \frac{\nu}{\nu - 1} > 1. \quad (121)$$

If supplier  $s$  activates exactly  $G$  products on link  $(s, b)$ , so that  $G_{sb} = G$  and the active products

---

<sup>8</sup>We maintain the constant-markup rule throughout. This is an analytical convenience. Under nested CES with  $\varepsilon \neq \nu$ , fully strategic multiproduct pricing would generally generate markups that depend on the inner-nest elasticity; see the discussion of multiproduct pricing in [Arkolakis et al. \(2021\)](#) and [Hottman et al. \(2016\)](#). The constant-markup approximation keeps variable profits proportional to revenues and allows us to isolate the scope-cannibalization mechanism on the demand side. It can be interpreted as a benchmark in which suppliers price product by product while taking the supplier-level price index as given.

have ranks  $g = 1, \dots, G$ , the within-supplier price index becomes

$$P_{sb}(G) = \tilde{\nu} \frac{w}{\phi_s} \left[ \sum_{g=1}^G h(g)^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}}. \quad (122)$$

Substituting (122) into (118) gives total link-level expenditure as a function of  $G$ :

$$R_{sb}(G) = E_b \left( \frac{P_b \phi_s}{\tilde{\nu} w} \right)^{\nu-1} \left[ \sum_{g=1}^G h(g)^{1-\varepsilon} \right]^{\frac{\nu-1}{\varepsilon-1}}. \quad (123)$$

Product-level revenue within the link, using the inner-nest allocation in (119), is

$$r_{sbg}(g; G) = R_{sb}(G) \frac{h(g)^{1-\varepsilon}}{\sum_{j=1}^G h(j)^{1-\varepsilon}}, \quad g = 1, \dots, G. \quad (124)$$

Equation (124) makes product-level cannibalization explicit. When the supplier adds a product, the denominator  $\sum_{j=1}^G h(j)^{1-\varepsilon}$  changes, reallocating revenue across all infra-marginal products on the link. Since  $\nu > 1$  and  $\varepsilon > 1$ , the exponent  $(\nu-1)/(\varepsilon-1)$  in (123) is positive, so expanding link-level scope raises total link revenue even though the revenue of individual infra-marginal products may fall.

Under the constant-markup rule, total variable profits are a constant fraction of link revenue:

$$\sum_{g=1}^G \pi_{sb}^V(g; G) = \frac{1}{\nu} \sum_{g=1}^G r_{sbg}(g; G) = \frac{1}{\nu} R_{sb}(G). \quad (125)$$

This is the key analytical simplification. Although product-level revenues are jointly determined through within-link cannibalization, total variable profits on the relationship depend on  $G$  only through the scalar  $R_{sb}(G)$ . The buyer-level problem therefore remains a one-dimensional choice over link-level scope.

### O.3 The supplier's two-stage problem under nested CES

We preserve the two-stage logic of the model. In Stage 1, supplier  $s$  chooses firm-level adopted scope  $G_s$  and pays the firm-level product-development cost  $F^P(G_s)$ . In Stage 2, for each potential buyer  $b$ , the supplier decides whether to activate the relationship, paying  $f_{sb}^L$ , and how many of its developed products to activate on that link, paying the corresponding product-activation costs. Relative to the deterministic-ladder problem in Appendix N, the key change is that the within-link revenue function  $R_{sb}(G)$  now incorporates cannibalization through the inner CES nest.

**Stage 2: link-level value function.** Given firm-level scope  $G_s$ , the maximum operating value that supplier  $s$  can obtain from buyer  $b$  is

$$V_{sb}(G_s; \phi_s) = \max_{G \in \{0, 1, \dots, G_s\}} \left\{ \frac{1}{\nu} R_{sb}(G) - \mathbf{1}\{G > 0\} f_{sb}^L - F_{sb}^A(G) \right\}, \quad (126)$$

or, substituting (123),

$$V_{sb}(G_s; \phi_s) = \max_{G \in \{0, 1, \dots, G_s\}} \left\{ \frac{E_b}{\nu} \left( \frac{P_b \phi_s}{\tilde{\nu} w} \right)^{\nu-1} \left[ \sum_{g=1}^G h(g)^{1-\varepsilon} \right]^{\frac{\nu-1}{\varepsilon-1}} - \mathbf{1}\{G > 0\} f_{sb}^L - F_{sb}^A(G) \right\}. \quad (127)$$

The first term is total variable profit from activating  $G$  products on the link, the second is the link fixed cost, paid only when the relationship is active, and the third is the cumulative product-activation cost. Optimal link-level scope and link activation are therefore characterized jointly by

$$G_{sb}^*(G_s; \phi_s) \in \arg \max_{G \in \{0, 1, \dots, G_s\}} \left\{ \frac{1}{\nu} R_{sb}(G) - \mathbf{1}\{G > 0\} f_{sb}^L - F_{sb}^A(G) \right\}, \quad (128)$$

with the relationship activated if and only if  $G_{sb}^*(G_s; \phi_s) \geq 1$ . The corresponding active buyer set at firm-level scope  $G_s$  is

$$\mathcal{B}_s(G_s) = \{b \in \mathcal{B} : G_{sb}^*(G_s; \phi_s) \geq 1\}. \quad (129)$$

Two features distinguish (128) from the deterministic-ladder Stage-2 problem in Appendix N. First, the gain from adding another product to the link is no longer additive rank by rank, because the inner-nest term

$$\left[ \sum_j h(j)^{1-\varepsilon} \right]^{(\nu-1)/(\varepsilon-1)}$$

depends on the entire set of active products. Second, adding a product changes not only the revenue generated by that product but also the revenues of infra-marginal products through the cannibalization channel. The closed-form rank cutoff  $\tilde{G}_{sb}(\phi_s)$  of the deterministic-ladder benchmark therefore no longer applies. Instead, the link-level problem must be solved as a discrete choice over  $G \in \{0, 1, \dots, G_s\}$ . Different buyers can optimally choose different scopes because  $E_b$ ,  $f_{sb}^L$ , and  $F_{sb}^A(\cdot)$  differ across relationships.

**Stage 1: firm-level scope.** Substituting the Stage-2 value functions into the supplier's outer problem gives

$$\Pi_s(G_s; \phi_s) = \sum_{b \in \mathcal{B}} V_{sb}(G_s; \phi_s) - F^P(G_s), \quad (130)$$

and optimal firm-level scope satisfies

$$G_s^*(\phi_s) \in \arg \max_{G_s \in \{0,1,\dots,G_{\max}\}} \Pi_s(G_s; \phi_s). \quad (131)$$

The outer logic is the same as in the baseline model of Section 4 and the deterministic-ladder benchmark of Appendix N. For each candidate  $G_s$ , the supplier aggregates the value of all potential buyer relationships and subtracts the firm-level product-development cost  $F^P(G_s)$ . The amortization mechanism is therefore preserved: the firm-level development cost is paid once, whereas the return to a developed product can accrue across multiple buyer relationships. A larger and more valuable buyer base consequently raises the return to expanding firm-level scope.

**Joint determination and solution.** Firm-level product scope and the active buyer set remain jointly determined. For any candidate  $G_s$ , equation (128) determines the optimal link-level decision for every potential buyer and therefore the active buyer set  $\mathcal{B}_s(G_s)$ . The supplier then evaluates the resulting value in (130) and chooses the maximizing  $G_s$ . Because the product-scope choice set is finite, the problem can be solved directly by evaluating Stage 2 for each candidate firm-level scope, in the same spirit as the exact discrete solution described in Section 4.4. The nested-CES extension therefore changes the within-link revenue calculation but does not create an additional equilibrium-existence problem.

**Comparison with the deterministic-ladder benchmark.** Conditional on  $G_s$ , each buyer-level problem in (127) remains a scalar discrete choice. What is lost relative to Appendix N is the simple closed-form rank-cutoff characterization of link-level scope. When  $\varepsilon = \nu$ , the inner aggregator in (123) reduces to

$$\sum_{g=1}^G h(g)^{-(\nu-1)},$$

and the additive deterministic-ladder structure is recovered. The nested-CES specification therefore generalizes the rank-based benchmark by introducing within-link cannibalization whenever  $\varepsilon \neq \nu$ , with  $\varepsilon > \nu$  corresponding to the empirically relevant case in which a supplier's own products are closer substitutes for one another than different suppliers' product lines are.

## O.4 Complementarity in both directions: buyers $\leftrightarrow$ products

The two-way interaction between buyer base and firm-level product scope established in Section 4.5 continues to operate under nested CES, although the within-link demand channel is modified by cannibalization.

**More buyers  $\Rightarrow$  more products.** Holding supplier productivity  $\phi_s$  fixed, expanding the active buyer set raises the return to firm-level product development because the cost  $F^P(G_s)$  is paid once and can be amortized across more relationships. Formally, the marginal gain from raising firm-level scope from  $g - 1$  to  $g$  is

$$\Delta\Pi_s(g; \phi_s) = \sum_{b \in \mathcal{B}} [V_{sb}(g; \phi_s) - V_{sb}(g - 1; \phi_s)] - f^P(g). \quad (132)$$

Each additional buyer can contribute a non-negative increment to the gross return from the marginal product, while  $f^P(g)$  remains fixed at the firm level. The same amortization logic emphasized in Section 4.5 therefore implies that, holding productivity fixed, optimal firm-level scope is weakly increasing in the size and value of the buyer base.

**More products  $\Rightarrow$  more buyers.** Holding  $\phi_s$  fixed, increasing  $G_s$  weakly expands the feasible set in the link-level problem (128). The maximized relationship value  $V_{sb}(G_s; \phi_s)$  is therefore weakly increasing in  $G_s$ . A larger firm-level portfolio allows a wider link-level product range on relationships for which the firm-level cap was binding. Through (118), a broader active product line lowers the inner-nest price index and can raise the supplier's share of the buyer's expenditure. Previously unprofitable buyer relationships can consequently become profitable as  $G_s$  rises. The active-set inclusion

$$\mathcal{B}_s(G_s) \subseteq \mathcal{B}_s(G_s + 1)$$

therefore continues to hold, paralleling the corresponding result in Section 4.5.

**Joint determination.** The two directions reinforce one another. More buyers raise the return to firm-level product development through amortization, while a broader product portfolio raises the value of potential buyer relationships. Nested CES changes how additional products affect revenues within each active link, but it does not eliminate the qualitative interaction between buyer base and product scope that motivates the model and the empirical co-movement documented in Fact 3.

## P Extension 2: Dynamic buyer accumulation over the life cycle

The baseline model in Section 4 and the nested-CES robustness exercise in Appendix O are static within-period environments. They characterize how a given buyer base affects firm-level product scope and how product scope, in turn, affects the profitability of buyer relationships, but

they do not explain how the buyer base itself is accumulated over time.

The life-cycle evidence in Appendix C.4 provides a natural motivation for adding this margin. Conditional on initial sales, younger suppliers rely relatively more on the joint addition of new buyers and new products, whereas the importance of selling established products to new buyers rises with supplier age. Because age is not a state variable in the benchmark model, these patterns are treated as complementary descriptive evidence rather than quantitative targets. Here we show how a simple buyer-accumulation block can generate a dynamic interpretation of those patterns.

## P.1 Buyer-base state variable

The static matching model's approximate scope rule in Section 4.5 depends on the match-adjusted buyer-base aggregator

$$\Lambda_s^M \equiv \sum_{b \in \mathcal{B}_s} D_b^{\theta_\zeta} (f_{sb,0}^A)^{1-\theta_\zeta}.$$

To introduce gradual buyer accumulation, let

$$\Lambda_{s,a}^M \geq 0 \tag{133}$$

denote the corresponding buyer-base state of supplier  $s$  at age  $a \in \{0, 1, 2, \dots\}$ . Define the contribution of buyer  $b$  to this aggregator as

$$\lambda_{sb} \equiv D_b^{\theta_\zeta} (f_{sb,0}^A)^{1-\theta_\zeta}.$$

Then

$$\Lambda_{s,a}^M = \sum_{b \in \mathcal{B}_{s,a}} \lambda_{sb} = N_{s,a} \bar{\lambda}_{s,a},$$

where  $N_{s,a} = |\mathcal{B}_{s,a}|$  is the number of active buyers and

$$\bar{\lambda}_{s,a} \equiv \frac{1}{N_{s,a}} \sum_{b \in \mathcal{B}_{s,a}} \lambda_{sb}$$

is their average contribution to the adoption aggregator. The state therefore summarizes both buyer base and the average value of those relationships for product-development incentives.

In each period, supplier  $s$  chooses buyer-acquisition effort

$$n_{s,a}^B \geq 0, \tag{134}$$

which determines the expected number of new buyer relationships reached at age  $a$ . Let  $m^B > 0$

denote acquisition or matching efficiency, let  $\bar{\lambda}^{\text{new}} > 0$  denote the average contribution of a newly acquired relationship to the buyer-base aggregator, and let  $\delta^B \in (0, 1)$  denote the per-period survival probability of an existing buyer relationship. The expected law of motion for the buyer-base state is

$$\mathbb{E} [\Lambda_{s,a+1}^M \mid \Lambda_{s,a}^M, n_{s,a}^B] = \delta^B \Lambda_{s,a}^M + m^B \bar{\lambda}^{\text{new}} n_{s,a}^B. \quad (135)$$

The buyer base next period therefore combines surviving current relationships with newly acquired buyers. We keep  $\delta^B$  age-invariant in this stylized extension so that life-cycle differences in buyer-base growth arise through acquisition rather than mechanically through an age-varying separation rate. This is consistent with the role of the survival evidence in Appendix B, while keeping the dynamic block deliberately parsimonious.

## P.2 Cost of buyer accumulation

Building the buyer base requires a convex networking or acquisition cost,

$$C_a^B(n_{s,a}^B) \equiv \xi_a^B \frac{f^B}{\gamma_B} (n_{s,a}^B)^{\gamma_B}, \quad \gamma_B > 1, \quad f^B > 0, \quad (136)$$

where  $\xi_a^B \geq 1$  is an optional age-specific networking wedge. The convexity parameter  $\gamma_B > 1$  makes rapid buyer acquisition increasingly costly and therefore gives firms an incentive to smooth the accumulation of customer relationships over time. If  $\xi_a^B$  is higher at young ages, young firms face an additional impediment to building their buyer base quickly, strengthening the gradual-accumulation mechanism.

The cost  $C_a^B$  can be interpreted as networking, sales, promotion, search, or prospecting effort aimed at forming new firm-to-firm relationships. The substantive assumption is simply that buyer acquisition is costly and that the marginal cost of rapidly expanding the customer base rises with acquisition effort.

## P.3 Dynamic supplier problem

The dynamic block is layered on top of the static matching problem in Section 4. Conditional on current firm-level scope and the current buyer-side state, product–buyer activation and link formation remain within-period decisions. In particular, the supplier’s link-formation decision follows the expected-surplus rule in equation (23), while realized products within an active relationship are selected according to equation (27).

To keep the dynamic extension tractable, we summarize the buyer-side state relevant for product-development incentives by  $\Lambda_{s,a}^M$ . Let

$$\Pi_s^{\text{flow}}(G_s; \phi_s, \Lambda_{s,a}^M) \quad (137)$$

denote the current-period expected profit block inherited from the static matching model, with the buyer-base state summarized by  $\Lambda_{s,a}^M$ . The supplier's dynamic problem is

$$W_s(\phi_s, a, \Lambda_{s,a}^M) = \max_{G_s, n_{s,a}^B} \left\{ \Pi_s^{\text{flow}}(G_s; \phi_s, \Lambda_{s,a}^M) - C_a^B(n_{s,a}^B) + \beta \mathbb{E}[W_s(\phi_s, a+1, \Lambda_{s,a+1}^M)] \right\}, \quad (138)$$

subject to the law of motion (135), where  $\beta \in (0, 1)$  is the discount factor. Each period, the firm therefore chooses both its current product scope  $G_s$ , inherited from the Stage-1 decision of the static model, and buyer-acquisition effort  $n_{s,a}^B$ . The firm internalizes that acquisition today raises the future buyer-base state and therefore changes the future return to product development.

**Static versus dynamic decisions.** Conditional on  $(G_s, \Lambda_{s,a}^M)$ , the product–buyer and link-level decisions remain within-period. The supplier evaluates potential relationships as in Sections 4.2–4.3: expected product surplus determines whether a link is formed, and realized product–buyer matches determine which developed products are activated within an active relationship. What the dynamic block adds is a decision above this static problem. The firm chooses  $n_{s,a}^B$  today to shape the future evolution of  $\Lambda_{s,a+1}^M$ , taking into account that a larger future buyer base raises the return to firm-level product scope through the amortization mechanism characterized in Section 4.5.

**Euler-style condition for buyer acquisition.** The first-order condition with respect to  $n_{s,a}^B$  takes the standard intertemporal form. Differentiating (138) and using (135),

$$\underbrace{\xi_a^B f^B(n_{s,a}^B)^{\gamma_B-1}}_{\text{marginal acquisition cost}} = \beta m^B \bar{\lambda}^{\text{new}} \cdot \underbrace{\mathbb{E} \left[ \frac{\partial W_s(\phi_s, a+1, \Lambda_{s,a+1}^M)}{\partial \Lambda_{s,a+1}^M} \right]}_{\text{marginal value of the buyer base}}. \quad (139)$$

The right-hand side links the dynamic acquisition decision to the static scope mechanism. An additional unit of the buyer-base state raises future profits directly and, through amortization, increases the future return to developing a broader product portfolio.

## P.4 Connection to the benchmark scope rule

Section 4.5 provides an approximate characterization of optimal firm-level scope in the interior Pareto region. Using the match-adjusted buyer-base aggregator, the static relationship is

$$G_s^* \propto \left[ \frac{\phi_s^{\theta_\zeta(\sigma-1)} \Lambda_s^M}{f^P(1)} \right]^{\frac{1}{\delta + \kappa \theta_\zeta - \mu}}, \quad \kappa = \mu + \alpha(\sigma - 1).$$

Embedding the same logic in the dynamic environment gives the life-cycle scope rule

$$G_s^* (\phi_s, \Lambda_{s,a}^M) \approx \text{integer} \left\{ \left[ \frac{\phi_s^{\theta_\zeta(\sigma-1)} \Lambda_{s,a}^M}{f^P(1)} \right]^{\frac{1}{\delta + \kappa\theta_\zeta - \mu}}, \right. \quad (140)$$

where the expression is rounded to the feasible integer product-scope choice. Holding productivity fixed, gradual accumulation of  $\Lambda_{s,a}^M$  therefore generates gradual expansion of optimal firm-level scope whenever the denominator

$$\delta + \kappa\theta_\zeta - \mu$$

is positive, as required by the finite-scope condition underlying the static approximation.

The same qualitative logic carries over to the nested-CES robustness exercise of Appendix O. A larger buyer base raises the payoff to paying the firm-level development cost  $F^P(G_s)$  and can also make additional buyer relationships profitable. Within-link cannibalization prevents the same closed-form characterization, but it does not overturn the basic monotonicity: a larger and more valuable buyer base raises the return to firm-level product development.

## P.5 Equivalent truncated-history representation

For comparison with dynamic production-network and customer-accumulation models such as Einav et al. (2021) and Aekka and Khanna (2024), the buyer-base state can equivalently be represented as a truncated history of past acquisition decisions. Iterating (135) forward gives

$$\Lambda_{s,a}^M \approx m^B \bar{\lambda}^{\text{new}} \sum_{\tau=0}^{A_B} (\delta^B)^\tau n_{s,a-1-\tau}^B, \quad (141)$$

where  $A_B$  is a truncation horizon chosen so that contributions from sufficiently old acquisition cohorts are negligible. The approximation is accurate when buyer relationships do not survive indefinitely, because the contribution of past cohorts decays geometrically at rate  $\delta^B$ . The dynamic buyer side can therefore be summarized either by the scalar state  $\Lambda_{s,a}^M$  in (135) or, equivalently, by a finite-dimensional history of recent buyer-acquisition decisions.

## P.6 Interpretation and mapping to the empirical facts

The static and dynamic blocks provide a unified interpretation of the empirical patterns. The baseline matching model in Section 4 explains why buyer base and product scope co-move through amortization and buyer-specific matching. The nested-CES extension in Appendix O shows how this interaction is modified when products within a supplier compete for expenditure inside a

buyer relationship. The present extension adds a dynamic margin by allowing the buyer base itself to be accumulated gradually over time.

This dynamic mechanism provides a natural interpretation of the life-cycle evidence documented in Fact 2 (Section 3.2) and Appendix C.4. Figure 19 shows that, conditional on initial sales, younger suppliers rely relatively more on the joint new-buyer–new-product margin (NLNP), whereas the importance of selling existing products to new buyers (NLOP) rises with supplier age. OLNP remains comparatively small across age groups.

In the dynamic model, young firms begin with a relatively small buyer-base state  $\Lambda_{s,a}^M$ . Because expanding that state is costly, buyer accumulation occurs gradually. Since optimal firm-level scope is increasing in  $\Lambda_{s,a}^M$  through (140), product scope expands alongside the buyer base. Early in the life cycle, increases in the buyer base therefore tend to coincide with product expansion, providing a natural counterpart to the relatively high NLNP share among younger suppliers. As  $\Lambda_{s,a}^M$  grows, the firm operates with a broader established product portfolio, so subsequent buyer acquisition is increasingly accommodated by extending existing products to new customers, consistent with the rising importance of NLOP among older suppliers in Figure 19.

The extension should be interpreted as a qualitative dynamic rationalization rather than as a calibrated model of the life cycle. A natural next step would be to estimate the acquisition-cost curvature  $\gamma_B$ , the age profile of  $\xi_a^B$ , and the persistence parameter  $\delta^B$  using the observed dynamics of buyer relationships, and then quantify the relative contributions of slow buyer accumulation and gradual product development to firm growth over the life cycle.