Supplier Search and Market Concentration *

Chek Yin Choi†

August 10, 2026

[Click Here to Newest Version]

Abstract

This paper studies how easier access to intermediate input suppliers affects productivity and market concentration. I develop a quantitative model of supplier search in which firms incur fixed costs to discover and bargain with input suppliers. The model provides a microfoundation for how input trade influences aggregate productivity and resource allocation. Evidence from firm-level import data motivates the framework through four patterns: growing dispersion in imported varieties, rising inequality in importer sales, lower input prices for larger firms, and stronger supplier network expansion in municipalities with better digital infrastructure. In general equilibrium, lower input search frictions reallocate resources toward more productive firms. Quantitatively, the mechanism accounts for about 27 percent of productivity growth and about 46 percent of the observed rise in concentration in the Swedish manufacturing sector from 1998 to 2021. A 10 percent tariff on imported inputs offsets the GDP gain and lowers concentration, providing a policy counterpart to the supplier search mechanism.

Keywords: Input markets; Firm-to-firm trade; Search costs; Market concentration; Tariffs

*I am very grateful to Timo Boppart for his guidance and support. I would like to also express my sincere gratitude to Michele Fioretti, Kurt Mitman, Per Krusell, Francis Kramarz, Josh Weiss, Horng Chern Wong, Amalia Repele, Mitch Downey, Brian Higgins, Kieran Larkin, Peter Nilsson, Yimei Zou, Juan Llavador, Gianmarco Ottaviano, Nevine El-Mallakh, Markus Kondziella, Luigi Iovino, Jan Bakker, Mikhail Oshmakashvili, Julian Hinz, Giuseppe Berlingieri, Nuriye Bilgin, and seminar participants at Stockholm University, Bocconi University, Kiel University and conference participants at NEWER Workshop, Swiss-Swedish Macroeconomics Workshop, 3rd RISE Workshop, SUDSWEC, 14th National PhD Workshop in Finance whose insightful comments significantly improved the quality of this paper. All mistakes are my own.

†IIES, Stockholm University, chek-yin.choi@iies.su.se, https://hkchekc.github.io/

1 Introduction

Market concentration has risen steadily since the 1990s in many advanced economies, sparking debate over its causes and consequences for productivity and allocative efficiency. Over the same period, it has become easier and cheaper for firms to source intermediate goods. Improvements in digital communication, transportation technology, together with growing sophistication of global supply networks and trade policies, have reduced search frictions and broadened the range of available suppliers. When input markets become more efficient, firms gain access to a larger and higher quality pool of suppliers, which lowers production costs and raises productivity (Amiti and Konings, 2007; Goldberg et al., 2010; Halpern et al., 2015). However, it is unclear if the gains are evenly distributed. This raises important questions: which firms benefit and by how much; what are the implications for aggregate productivity; and how do these changes influence firm entry into importing. In this paper, I argue that rising productivity and increasing market concentration are two sides of the same coin, both driven by improvements in input market efficiency.

I begin my analysis by establishing three key empirical patterns using detailed Swedish administrative data in the manufacturing sector. These facts provide new systematic evidence linking firm-level sourcing behavior to rising firm concentration. First, there is substantial dispersion in the number of imported varieties across firms. Further, firms at the upper tail now source about 60% more inputs than twenty years ago, making the dispersion even larger. Second, the sales distribution among importing firms has become considerably more unequal: the largest importers are 75–125% larger than they were in 1998, while the smallest importers are about 50% smaller. Third, firms that are 10% larger pay about 1.6% lower unit values.

I rationalize these empirical findings through the lens of a model of monopolistic competition with frictional input markets. The model incorporates the necessary ingredients to explain the facts documented above and to study how frictions in input markets affect allocative efficiency as well as market concentration. Specifically, I embed a random search and bargaining framework into a model of two-sided heterogeneity and monopolistic competition. Buyers (downstream firms) can produce in-house or source intermediate inputs from sellers (upstream firms). Upstream firms are heterogeneous in their productivity and produce with only labor. Downstream firms are also heterogeneous in their productivity and produce using a Cobb-Douglas aggregator between labor and a CES aggregator of intermediate goods that features love of variety.

Downstream firms can only source intermediate inputs from upstream firms in their supplier network. Buyers can enlarge their supplier network by sequentially searching for new sellers. A matched buyer and seller then bargain over quantities and prices for the intermediate good. The firm size distribution and supplier-network size distribution are thus endogenously determined in equilibrium.

The key mechanism of the model is that there is endogenous selection into both the extensive margin (the decision to start buying from an upstream firm) and the intensive margin (the number of upstream firms a buyer chooses to buy from). The mechanism allows the model to match the cross-sectional facts: more productive firms search more, build larger supplier networks, become larger, and pay lower unit prices due to their better outside options in bargaining. A reduction in search frictions allows productive firms to expand their networks and reduce the prices of intermediate inputs further, reallocating resources toward them. Aggregate efficiency rises as production shifts toward more productive firms and as broader supplier networks enhance input complementarities. The firm size distribution widens, and market concentration increases.

To quantify these mechanisms, I estimate the model to match the key moments documented in the Swedish manufacturing data. The key endogenous parameters in the estimation are the dispersion in firm productivity, which is estimated to reproduce the observed top sales share, and the decline in search costs, which is estimated to match the observed rise in the mean number of imported varieties. In the model, search frictions must fall by about 40 percent to match the observed increase in imported varieties. The decline in frictions leads to a 13 percent increase in real GDP and an increase in the Herfindahl-Hirschman index (HHI) by about 12.7 percent. For comparison, Sweden's real manufacturing value added per employee increased by approximately 49 percent over 1998 to 2021 (SCB, 2025a,b,c,d), and market concentration rose by about 28 percent over the same period. These figures imply that the mechanism accounts for roughly 27 percent of observed labor productivity growth and about 46 percent of the increase in concentration.

To provide further empirical evidence for the mechanism linking input market efficiency and firm concentration, I use the staggered rollout of the high-speed broadband infrastructure across Swedish municipalities as plausibly exogenous variation in input search costs. I show that firms located in municipalities with lower search costs, proxied by better high-speed internet coverage, expand their supplier networks and particularly among the largest firms. High-speed internet coverage leads firms to import 5 more varieties on average, with the largest firms importing around 23 more varieties.

I next compare the decline in search frictions with an alternative counterfactual in which foreign suppliers become more efficient over time. In this exercise, I vary only the mean productivity of foreign suppliers while keeping search costs fixed. The magnitude of the increase in productivity is disciplined to match the decrease in the demand shifter in the main exercise, which is a composite of aggregate price and production. Rising supplier efficiency delivers a similar increase in aggregate productivity and market concentration but mainly through the intensive margin. It does not match the expansion in the number of imported varieties observed in the data, while lower search costs do. This pattern suggests that both forces could have contributed, with declining search frictions playing a central role in driving the observed rise in market concentration.

Finally, I use the model to study a timely and topical policy question: the impact of a uniform tariff on imported intermediates. This exercise links technological changes in input market efficiency to tangible policy instruments that affect firms in comparable ways. Although tariffs are not identical to search costs, both raise the effective cost of accessing foreign inputs and discourage the formation of supplier relationships. As a result, a tariff compresses supplier networks and shifts activity toward smaller and less productive firms, replicating the economic effects of higher search frictions. Quantitatively, a 10 percent tariff offsets the output gains and the increase in market concentration generated by the decline in search costs over the past two decades. While such a policy lowers aggregate productivity, it narrows performance differences between large and small firms, which can make it politically attractive.[1]

This paper contributes to the literature in three ways. First, I examine the role of changes in input markets as a determinant of market concentration in downstream industries. In doing so, I complement existing studies that emphasize other drivers of concentration, such as rising entry barriers (Covarrubias et al., 2020), intensified import competition (Amiti and Heise, 2025), declining market-spanning costs due to advances in information technology (Aghion et al., 2023), and the rise of intangible capital (Crouzet and Eberly, 2019; De Ridder, 2024; Weiss, 2020; Bajgar et al., 2025).

The literature has studied how changes in trade patterns can reshape the firm size distribution in profound ways. For example, Amiti and Heise (2025) show that intensified import competition raises market concentration among local firms by displacing less productive domestic producers. This perspective, however, focuses on trade in final goods, while the role of intermediate inputs in shaping firm dynamics has received much less attention. Imported final goods directly compete with domestic producers and reduce demand for their output, a demand shock. By contrast, importing inputs

[1] Since large and highly productive firms are concentrated in major urban regions, restricting input trade slows their expansion and limits further spatial concentration of economic activity, a recurrent theme in European policy discussions (OECD, 2023; European Commission Directorate General for Regional and Urban Policy, 2024).

generates a productivity shock for downstream firms by lowering effective production costs. This distinction shifts the focus from how trade reallocates demand across firms (Eaton et al., 2011) to how heterogeneity in access to international suppliers influences the production efficiency of firms. Understanding this margin is particularly important because intermediate inputs account for roughly 56 percent of global trade (Miroudot et al., 2009) and dominate the EU's import basket. Yet we know relatively little about how improvements in input market efficiency translate into firm-level outcomes.

Second, this paper contributes to the growing firm-to-firm trade literature, which studies how individual buyer and supplier relationships shape trade patterns, prices, and productivity. Using detailed Swedish import data, I document substantial heterogeneity in firms' sourcing behavior and input prices, consistent with Atalay (2014). Prior studies attribute such patterns to buyer market power (Morlacco, 2020; Rubens, 2023), input quality differences (Kugler and Verhoogen, 2012), and match-specific frictions (Burstein et al., 2024). In Indian firm-to-firm data, Aekka et al. (2026) find that younger firms have fewer suppliers and face higher input costs, and attribute this to costly search for trading partners. These patterns motivate my theoretical framework that explains why more productive and more connected firms obtain systematically lower input prices.

I develop a model of buyer and supplier relationships that brings together the key ingredients emphasized in this literature, including bargaining (Alviarez et al., 2023; Grossman et al., 2023), search (Eaton et al., 2022a), and matching (Eaton et al., 2022b). The model allows for endogenous relationship formation and bilateral bargaining over both quantities and prices, and is related to Oberfield (2018), who studies how production linkages shape aggregate productivity. The framework contributes to general search and bargaining theory by combining explicit search costs, two-sided heterogeneity, and multidimensional bargaining in a single model. To my knowledge, existing work incorporates at most two of these elements. The closest model is Grossman et al. (2024), who combine costly search with negotiated input prices, but under Leontief aggregation, homogeneous buyers, and one-dimensional bargaining. Canonical Diamond–Mortensen–Pissarides (DMP) search-and-matching models include explicit search costs but feature one-sided heterogeneity and one-dimensional bargaining (Pissarides, 2017). Extensions with two-sided heterogeneity in a DMP environment (e.g., Lise et al., 2016; Shimer and Smith, 2000) are also with only one-dimensional bargaining. Models with explicit search costs in two-sided settings exist (e.g., Eeckhout and Kircher, 2010; Eaton et al., 2022a), but abstract from bargaining.

Within this broader firm-to-firm framework, the paper also connects to the literature on imported intermediates, which shows that access to foreign inputs can raise firm productivity by lowering costs or improving input quality (Amiti and Konings, 2007; Goldberg et al., 2010; Halpern et al., 2015; Gopinath and Neiman, 2014). This paper complements that work by modeling how input market efficiency, shaped by search frictions, supplier heterogeneity, and bilateral bargaining, endogenously determines firms' access to imported inputs and the prices they face. In doing so, it links firm-level sourcing behavior to aggregate outcomes such as productivity and market concentration.

Third, I show that importing intermediate goods can generate distortions in the form of input price dispersion across firms. In contrast to standard misallocation models, where such wedges are exogenous, my framework derives them from input market frictions and bargaining. As search frictions decline, aggregate productivity increases and the market becomes more efficient, yet dispersion in both input prices and TFPR also rises because more productive firms gain better access to high-quality or low-cost suppliers. This mechanism complements the findings of Boehm and Oberfield (2020), who link input trade to allocative efficiency, and connects to the misallocation frameworks of Restuccia and Rogerson (2008) and Hsieh and Klenow (2009). It also relates to Edmond et al. (2015) and Dhingra and Morrow (2019), who study the procompetitive and welfare-improving effects of trade in consumer goods, by showing that input trade can raise aggregate productivity while reshaping the distribution of firm-level distortions.

The remainder of the paper is organized as follows. Section 2 describes the data and Section 3 documents the empirical facts. Section 4 develops the model, with two simplified cases in Sections 4.4 and 4.5. Section 5 maps model predictions to the data and quantifies the impact of falling search costs on productivity and concentration. Section 6 tests the mechanism against the staggered rollout of fiber broadband. Section 7 examines rising foreign supplier efficiency and compares it with the baseline mechanism. Section 8 presents the tariff counterfactual. Finally, Section 9 concludes.

2 Data

I combine three high-quality Swedish registry sources into a firm-level panel of manufacturing firms covering 1998 to 2021: detailed customs records of import transactions, firm financial statements, and municipality-level data on fiber internet diffusion. I merge them at the firm–year level and attach municipality characteristics by firm location. The panel contains 36,483 importing firms, 8,396 in an average year, sourcing 56,608 distinct varieties from 179 countries.

2.1 Import Data

The import dataset comes from the import component of the *Utrikeshandel med varor* (Foreign Trade in Goods) dataset compiled by Statistics Sweden (SCB). These import records report, for each firm-year, the country of shipment, the HS8 product code, the import value, and the physical quantity (usually weight). For some products, the dataset also includes an additional variable, "Other Quantities" (for example, number of pieces for pencils or m^2 for curtains).[2] I define a variety as an eight-digit product code by country, which corresponds to a supplier variety in the model.

The rich data allows me to construct unit values, which I interpret as prices. For each firm f, variety v, and year t, the unit price is the import value divided by the physical quantity:[3]

$$P_{f,v,t} = \frac{\text{Value}_{f,v,t}}{\text{Quantity}_{f,v,t}}.$$

To remove variation across products and countries, I compute the within-variety–year relative (residualized) price:

$$p_{f,v,t} = \log\left(\frac{P_{f,v,t}}{\bar{P}_{v,t}}\right),$$

where the benchmark price $\bar{P}_{v,t}$ is the value-weighted average price for variety v in year t:

$$\bar{P}_{v,t} = \frac{\sum_f \text{Value}_{f,v,t}}{\sum_f \text{Quantity}_{f,v,t}}.$$

I also construct firm-level indices for input price and quantity. The input price index measures the average log deviation of firm-specific prices from variety-level averages, weighted by import value:

$$\text{InputPrice}_{f,t} = \frac{\sum_v \text{Value}_{f,v,t}\, p_{f,v,t}}{\sum_v \text{Value}_{f,v,t}}.$$

The input quantity index $\text{InputQuantity}_{f,t}$ is defined analogously, replacing $p_{f,v,t}$ with $\log\left(\frac{\text{Quantity}_{f,v,t}}{\bar{\text{Quantity}}_{v,t}}\right)$.

The dataset includes all imports originating from outside the European Union (including Switzerland and EEA countries) and intra-EU imports for firms whose total annual imports exceed a reporting threshold. This threshold has gradually increased over time,

[2] A list of such products is available in the section *Varukoder (KN) för uppgiftslämnande till Intrastat* on the SCB website.

[3] Import values are deflated by the aggregate CPI.

from about SEK 1.5 million in the early 2000s to roughly SEK 9 million in recent years (see SCB (2018) for details).[4]

For data cleaning, I exclude observations with total import value below 100 SEK (approximately 9 USD) and those where the unit price is more or less than 15 times the average variety price, as these likely reflect non–arm's-length transactions, placeholder values, or other errors.

A limitation of the data is that if a firm imports the same variety from multiple suppliers within a year, the dataset reports only one aggregated record.

2.2 Balance Sheet Data

To characterize importing firms, I link the import data to firms' balance sheets using a common firm identifier available in both datasets. I use the industry codes (SNI) reported in the balance sheet data to restrict the sample to manufacturing firms. This linkage also allows me to compare the characteristics of importing firms with those of the broader firm population, as documented in Summary Statistic.

From the balance sheet data, I obtain firm-level measures of size, input expenditures, and productivity. I focus primarily on sales and employment and I construct measures of productivity, which I will explain in Productivity Measurement. I deflate all monetary variables, such as sales, using the aggregate CPI.

2.3 Internet Data

As a proxy for search costs, I use municipality-level internet connection data from the Swedish Post and Telecom Authority (PTS).[5] From these data I construct the share of firms in each municipality with access to fiber broadband, which I interpret as a proxy for lower search costs.[6] Each establishment in the administrative data is identified by a municipality code, which allows me to merge the internet coverage data to firms' geographic locations. The variable is available from 2010 to 2023, so the fiber analysis covers 2010 to 2021.

[4]Intra-EU trade is subject to this reporting threshold. When the United Kingdom left the EU in 2021, imports below the threshold from the UK that had previously been excluded began to appear in the data. This results in a mechanical increase in the number of small import records in 2021. To address this issue, I exclude post-2020 UK observations from trend plots.

[5]`https://statistik.pts.se/telekom-och-bredband/mobiltackning-och-bredband/dokument/`

[6]The dataset reports annual coverage by technology, including the total number of firms in each municipality and the number with access to each connection type.

3 Stylized Facts

This section documents three dimensions in which importing firms are very different from one another, and increasingly so: the number of varieties they import, their sales, and the input prices they pay. These facts motivate the main mechanism of the theory developed below, namely how search costs affect firms and, through them, the aggregate economy.

Search costs have declined substantially over the past three decades due to rapid improvements in transport and communication technologies. In Sweden, for instance, only about 27% of commercial establishments were connected to the internet through fiber in 2010, whereas by 2021 this share had increased to 79%. Lower search frictions reduce the barrier to entering import markets for small firms while allowing large incumbents to extend their supplier networks, and I read these facts through that lens.

3.1 Number of Imported Varieties

Importing firms differ widely in how many varieties they source. In an average year the median importer buys 4 varieties while the 95th percentile buys 85, with a mean of 20.8 and a standard deviation of 78.7. To illustrate changes in sourcing patterns, I rank firms within each year by the number of varieties they import and track selected percentiles of that distribution over time. For comparability, I normalize each percentile by its 1998 value.

I find that, except for the top percentiles, most firms have experienced little change in the number of imported varieties. In contrast, firms in the upper part of the distribution have expanded their variety of imports by 20% to up to 60% for the 99.9th percentile firms. Because the number of imported varieties reflects firms' search activities, this pattern suggests that lower search costs have had differential effects across the distribution. The same pattern holds when firms are ranked by sales instead, as shown in Number of varieties by sales percentile.

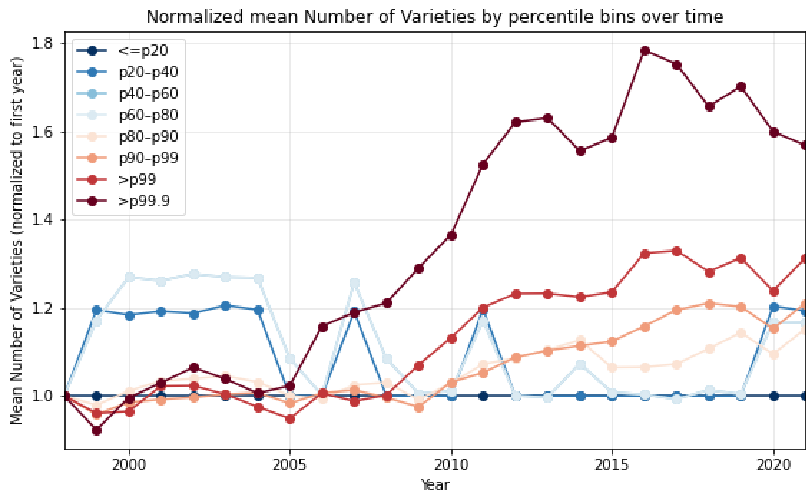

Figure 1: Changes in the number of imported varieties across variety percentiles

A large part of this pattern is driven by the entry of smaller firms into importing. When I restrict the sample to a balanced panel of firms that import every year in the data, all firms show an increase in the number of imported varieties relative to 1998, as shown in Number of varieties in a balanced panel.

The same pattern emerges in Figure 2, which reports the number of countries from which each firm sources a given 8-digit product (Panel A) and the number of distinct 8-digit products sourced by each firm in a given year (Panel B). In both cases, firms at the top of the distribution expand their sourcing scope more rapidly than others. The two panels decompose the growth in imported varieties into a country margin and a product margin, and neither dominates: firms at the top expand along both. This motivates a framework in which frictions operate supplier by supplier, rather than one in which the country is the level at which frictions arise, as in Antràs et al. (2017).

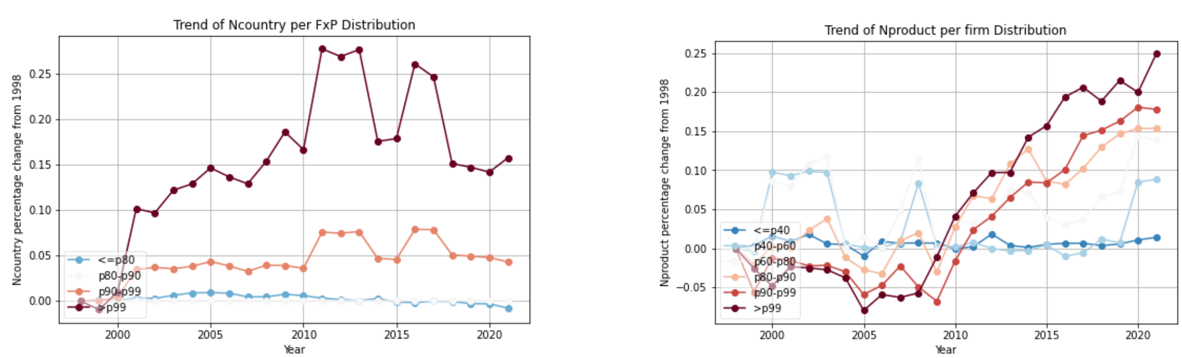

(a) Number of source countries per firm–product pair.

(b) Number of products sourced per firm.

Figure 2: Sourcing scope across firms, normalized to 1998.

3.2 Sales

Importers also differ sharply in size, and the gap has widened. I rank firms within each year by annual sales and track selected percentiles over time, normalizing each percentile by its 1998 value.

Sales fall substantially at the median and below: the 20th percentile firm in 2021 is approximately 50% smaller than its 1998 counterpart. This decline reflects entry, as search costs fell, less productive firms that had faced prohibitive barriers began to import. At the upper end, sales grew by up to 75% at the 99th percentile and 125% at the 99.9th. Among firms that were already importers in 1998, sales increased across the entire distribution (Sales dispersion in a balanced panel), so incumbents benefited as well.

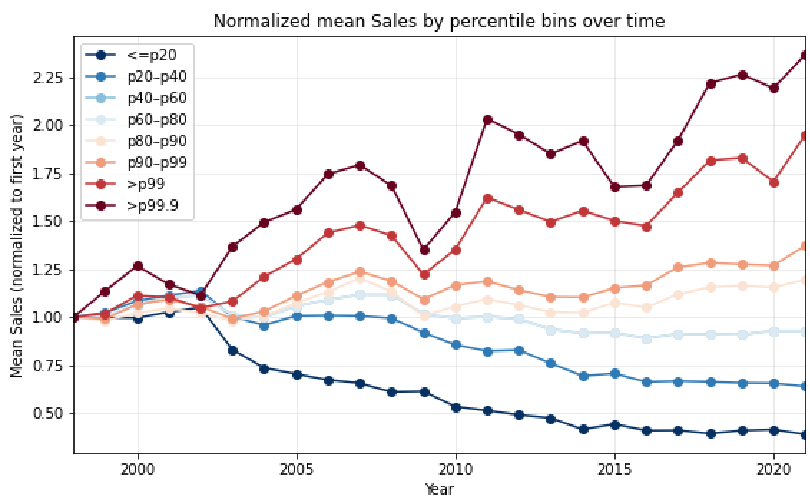

Figure 3: Changes in importing firm size across percentiles, 1998–2021

This fact is again consistent with falling search costs widening the size distribution along both the extensive margin, by lowering entry barriers for small firms, and the intensive margin, by reducing expansion costs for large ones.

3.3 Input Prices

Importers also differ in what they pay for comparable inputs. Input prices are another channel through which input markets affect productivity and allocative efficiency.

I relate input prices to firm size by estimating

$$\text{InputPrice}_{f,t} = \beta_0 + \beta_1 \log\left(\frac{\text{Size}_{f,t}}{\text{Size}_{i,t}}\right) + \epsilon_{f,t}, \tag{1}$$

where $\text{InputPrice}_{f,t}$ is the firm-level price index defined in Import Data and $\text{Size}_{f,t}$ is either the number of employees or the value-weighted average of import quantities, each demeaned by its industry–year average so that β_1 is an industry–year–relative elasticity. Estimates using other size measures, such as sales or TFPQ, are reported in Alternative measures and other samples.

Table 1: Regression Results: Input Prices and Firm Size

	(1) Employment	(2) Import Quantity
log No. of Workers	-0.1629^{***}	–
	(0.0017)	
Input Quantity	–	-0.2774^{***}
		(0.0014)
Observations	172,602	172,602
R^2	0.0534	0.2837

Notes: Dependent variable is $\text{InputPrice}_{f,t}$. Standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

I find a strong negative relationship: a 10% larger firm pays about 1.6% lower input prices when size is proxied by employment and about 2.8% lower when proxied by import quantities. This is consistent with buyer market power or bargaining power, whereby firms with better outside options obtain better terms.

Input prices have also become more dispersed. I measure dispersion as the variance of the residualized price $p_{f,v,t}$ defined in Import Data, so that variation across products and source countries is already netted out. Figure 4 shows that this variance rises steadily over the period, and the pattern is robust to restricting the sample to intra-EU imports or to many of Sweden's major trading partners.

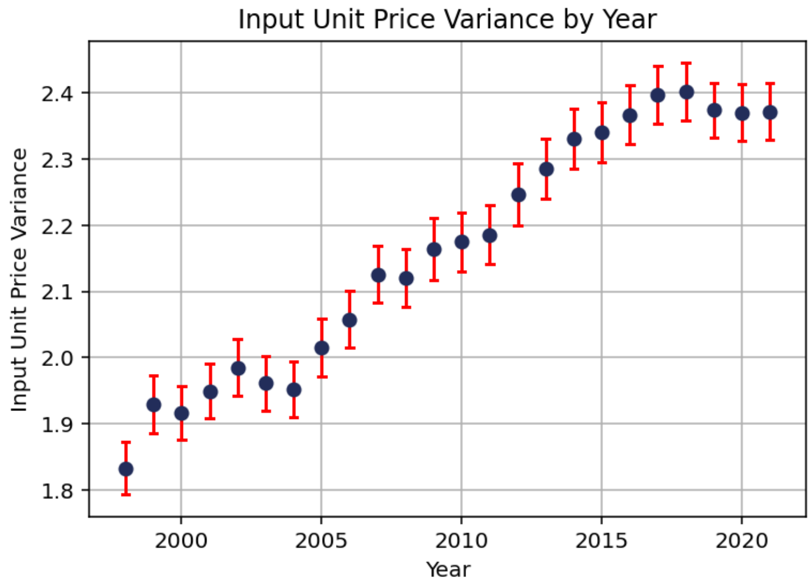

Figure 4: Input price dispersion over time

Most of this dispersion sits within firms rather than across them. To separate the two, I estimate

$$p_{f,v,t} = \alpha_{f,t} + \varepsilon_{f,v,t},$$

separately in each year, where the firm fixed effect $\alpha_{f,t}$ captures the component of the price a firm pays that is common to all of its varieties, and $\varepsilon_{f,v,t}$ the variation across varieties within the firm. Figure 5 splits the variance of $p_{f,v,t}$ into these two parts.[7] The across-firm component is roughly 0.5 in every year. The same buyer obtains one input cheaply and another expensively.

[7] The total here is higher than in Figure 4, which averages the variance of $p_{f,v,t}$ within each variety and so discards variation between varieties. The decomposition uses the pooled variance, which retains it. Both trend upward.

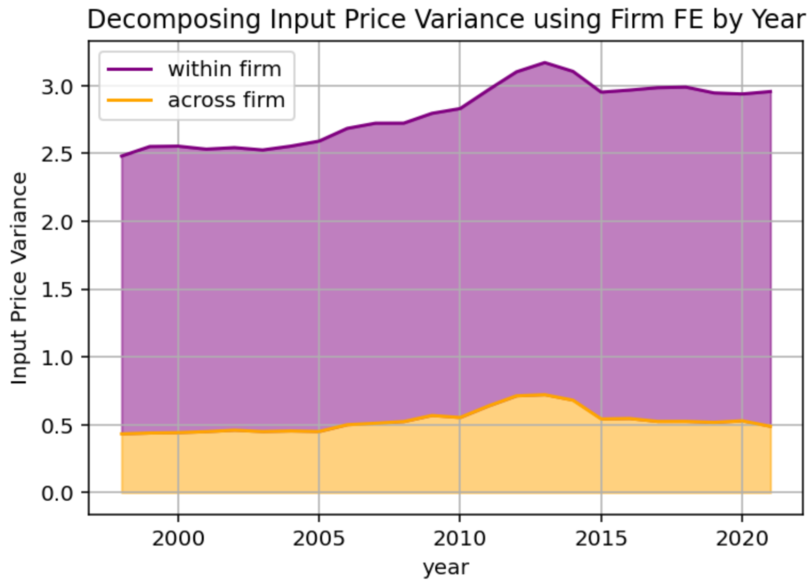

Figure 5: Within- and across-firm decomposition of input price variance

This is hard to reconcile with theories in which each firm carries a single markup or markdown, but it is what search operating supplier by supplier implies. A firm meets its suppliers one at a time, and the price it negotiates depends both on how efficient that particular supplier is and on how large its network already was when the contract was struck, so prices vary across the portfolio of a single buyer. As search costs fall and firms accumulate more suppliers, this within-firm spread widens alongside the smaller across-firm gap that the size gradient above describes.

Quality differences are unlikely to explain it. I compare price variance within differentiated and non-differentiated products, as classified by Rauch (1999). The measure is conservative, since the gap likely also reflects bargaining and information frictions, but neither the quality gap nor the share of differentiated products in total imports has increased over the period (Quality Differences).

3.4 Summary

To summarize, importing firms differ along three dimensions, and along each the differences have widened.

First, the distribution of imported varieties has become more unequal: firms at the top now import many more varieties than two decades ago, while firms at the bottom import roughly the same number as before.

Second, the sales distribution among importers has widened in the same way, linking changes in input sourcing to the growing gap in firm size.

Third, larger firms pay lower unit prices for comparable inputs, and the dispersion of those prices has risen. Firms that face lower input prices also earn higher profit shares (Input Price and Profit), tying input cost heterogeneity to firm performance.

Together these patterns suggest that falling search frictions allow large firms to expand their supplier networks more rapidly, raising productivity at the top and contributing to concentration. They motivate the key features of the model: the connection between supplier networks and firm scale, bargaining over input prices, and heterogeneous responses to lower search costs.

4 Theory

Motivated by these facts, I develop a model of buyer–supplier search and bargaining that links reductions in search frictions to supplier network formation and market concentration. The model builds on a standard framework with monopolistic competition in the final goods market, in which downstream buyers produce using labor and a CES composite of differentiated intermediate inputs. The novel contribution lies in the input procurement process: buyers must search to find upstream suppliers, and each match is settled by bilateral bargaining over both the quantity delivered and the payment. This structure captures how search frictions and contractual bargaining jointly shape the three dimensions documented above, namely the number of varieties a firm imports, the prices it pays, and ultimately its size.

4.1 Environment

The economy consists of upstream suppliers and downstream producers. Each upstream supplier u uses labor l to produce an intermediate input x_u according to a linear production technology:

$$x_u = z_u l,$$

where z_u denotes the supplier's productivity, so that the unit cost of production is w_{foreign}/z_u and w_{foreign} is the foreign wage. Without loss of generality I set $w_{\text{foreign}} = 1$, so supplier efficiency is fully captured by z_u.[8] The buyer observes z_u once a match is formed.

Downstream firms combine differentiated intermediate inputs with labor to produce the final good. The production function of a representative downstream firm is given by

$$y = zX^{\alpha}l^{1-\alpha}, \qquad X = \left(\sum_u x_u^{\rho}\right)^{\frac{1}{\rho}},$$

where y is output, z is the downstream firm's productivity, l is production labor, and X is the CES composite of all inputs sourced from upstream suppliers. I interpret love of variety in production as follows: adding a new variety provides a more specialized input, thereby increasing production efficiency, in a spirit similar to (Ethier, 1982). The parameter $\alpha \in (0, 1)$ is the expenditure share on intermediate inputs, while $\rho \in (0, 1]$ determines the elasticity of substitution across input varieties: a higher ρ makes inputs closer substitutes, a lower ρ greater complements.

The model timing is as follows: all downstream firms begin the period without any inputs. Each firm first decides whether to pay a search cost to look for suppliers or to stop searching altogether. A firm that chooses to search is randomly matched with a potential supplier. Upon meeting, the two parties bargain over a binding contract that specifies both the quantity of inputs to be delivered and the corresponding payment. After reaching an agreement, the downstream firm decides whether to terminate its search or to incur the search cost again to potentially find additional suppliers.

When all firms have completed their search and chosen to stop, production of the final good takes place simultaneously across all firms. The entire sequence of search, matching, bargaining, and production occurs instantaneously within the model period. To characterize equilibrium behavior, I solve the model by backward induction, starting from the production stage.

Two regularities from the literature support this timing. Buyer–supplier relationships are short-lived, typically lasting about one year (Martin et al., 2023), which justifies modeling the supplier network as re-optimized each period. Inventory considerations are minor at annual frequency, since importers restock roughly every 150 days on average (Alessandria et al., 2010).

[8] It would be the same if I instead defined an efficiency variable $e_u = w_{\text{foreign}}/z_u$, but since I do not distinguish between foreign countries that does not matter.

4.2 Search, Matching, and Bargaining

Consider a downstream firm with productivity z that has already completed its search and is matched with N upstream suppliers indexed by $u = 1, \ldots, N$. Let $\mathbf{x} = (x_1, x_2, \ldots, x_N)$ denote the vector of input quantities purchased from these suppliers. The composite input used in final production is

$$X(\mathbf{x}) = \left(\sum_{u=1}^{N} x_u^{\rho} \right)^{\frac{1}{\rho}}.$$

Given the input bundle $\mathbf{x}$, the firm chooses price p and labor l to maximize profits. Its value of operating without further search, $V^{NS}(z, \mathbf{x})$, is

$$V^{NS}(z, \mathbf{x}) = \max_{p,\, l \geq 0} \Big[p\, z\, X(\mathbf{x})^{\alpha} l^{1-\alpha} - wl \Big]$$

subject to product-market demand

$$y = C \left(\frac{p}{P} \right)^{-\epsilon},$$

where w is the wage rate, C is aggregate consumption, P is the aggregate price index, and $\epsilon > 1$ is the elasticity of substitution across differentiated final goods. This is the familiar monopolistic competition demand curve. This stage captures the production and pricing decisions of a firm that has concluded its supplier search and operates with its current network of relationships.

This problem admits a closed form solution expressing V^{NS} in terms of the state $(z, \mathbf{x})$ and parameters alone; see Closed Form Solution for Important Variables for the derivation.

Before stopping, a firm with productivity z and existing portfolio $\mathbf{x}$ decides whether to search for another supplier or to proceed directly to production. Its value function is

$$V(z, \mathbf{x}) = \max \Big\{ V^{S}(z, \mathbf{x}),\ V^{NS}(z, \mathbf{x}) \Big\},$$

where $V^{S}(z, \mathbf{x})$ is the value of searching for another supplier and $V^{NS}(z, \mathbf{x})$ the value of stopping, defined above.

If the firm chooses to search, it incurs a search cost κ measured in labor units and faces random matching with potential upstream suppliers. The expected value of searching is

$$V^{S}(z, \mathbf{x}) = \int V^{m}(z, z_u, \mathbf{x})\, dF(z_u) \;-\; w\, \kappa,$$

where z_u denotes the productivity of a prospective supplier drawn from the distribution F. The integral is the expected value of a match, taken before the supplier's productivity is realized. After paying the search cost, the firm proceeds to the bargaining stage, where it negotiates a binding contract specifying the input quantity x_{new} and the corresponding transfer payment T.

When a downstream firm with productivity z and existing input bundle $\mathbf{x}$ meets a new upstream supplier with productivity z_u, the two parties negotiate a binding contract over the quantity of the new input x_{new} and the transfer payment T. The outcome of this bilateral negotiation maximizes a Nash product of buyer and seller surpluses:

$$\max_{T,\, x_{\text{new}} \geq 0} \Big(V(z, \mathbf{x}_{\text{new}}) - T - V(z, \mathbf{x}) \Big)^{\theta} \Big(T - \tfrac{w_{\text{foreign}}}{z_u}\, x_{\text{new}} \Big)^{1-\theta},$$

where $\theta \in (0,1)$ is the bargaining weight of the buyer, $V(z, \mathbf{x}_{\text{new}})$ is the downstream firm's value after adding the new input, and $V(z, \mathbf{x})$ is its pre-match value. The transfer T is the payment from the buyer to the supplier, and $\frac{w_{\text{foreign}}}{z_u}\, x_{\text{new}}$ is the supplier's production cost.

The first term in the Nash product is the buyer's surplus from obtaining the new input, the second the seller's surplus from supplying it. The negotiated outcome determines both the optimal quantity x_{new} traded and the transfer T. The value of a successful match to the downstream firm is then

$$V^m(z, z_u, \mathbf{x}) = V(z, \mathbf{x}_{\text{new}}) - T,$$

which represents the firm's continuation value after accounting for the payment to the new supplier.

4.3 State Compression

Each of the value functions above takes the pair $(z, \mathbf{x})$ as its argument, and $\mathbf{x}$ grows by one element with every match the firm accepts. If each upstream supplier corresponds to one variety, the state has dimension $N+1$, which exceeds 1,000 for the largest importers in the data. Any dynamic problem in which an agent accumulates a variable number of heterogeneous objects has this feature. Here the CES structure removes it.

Proposition 1. *$V(z, \mathbf{x})$ depends on $\mathbf{x}$ only through the composite X, so $V(z, \mathbf{x}) \equiv V(z, X)$. The negotiated pair (x_{new}, T) inherits this, so the buyer's whole problem runs on the state (z, X).*

Proof. Output depends on inputs only through X, so V^{NS} does too. Adding a supplier expands the vector to $\mathbf{x}_{\text{new}} = (x_1, \ldots, x_N, x_{\text{new}})$, and the composite updates as

$$X_{\text{new}} = \left(\sum_{u=1}^{N} x_u^{\rho} + x_{\text{new}}^{\rho} \right)^{1/\rho} = \left(X^{\rho} + x_{\text{new}}^{\rho} \right)^{1/\rho},$$

so X_{new} is built from X and x_{new} alone. Suppose now that V depends on the portfolio only through X. The buyer's surplus in the Nash program, $V(z, \mathbf{x}_{\text{new}}) - T - V(z, \mathbf{x})$, is then a function of X, x_{new} and T; the solution (x_{new}, T) and the match value V^m therefore carry no more information about the portfolio than X does, and neither do $V^S = \int V^m \, dF - w\kappa$ or $V = \max\{V^S, V^{NS}\}$. The Bellman operator thus maps functions of (z, X) into functions of (z, X), and since this set is closed its fixed point lies in it. $\square$

The state is now two numbers rather than a list that grows without bound, and the reduction is exact.

Nothing in the argument is specific to bargaining, or to suppliers. The compression works whenever payoffs see the accumulated set only through the aggregate, and a new item enters through the aggregate and its own quantity: a consumer adding varieties to a CES basket, a firm adding product lines, a producer combining heterogeneous tasks.

Other methods compress variable-size states more generally, among them inclusive value sufficiency in industrial organization (Gowrisankaran and Rysman, 2012) and the moment representation of permutation-invariant payoffs (Fukasawa, 2025). Proposition 1 trades that generality for convenience: one line of algebra, and CES is already a workhorse utility and production function in structural work.

This technique supports the rich framework of this paper, which incorporates search, two-sided heterogeneity and bargaining over both prices and quantities. Firms end up with portfolios of their own, and the state stays two-dimensional regardless. That is what allows a single model to account for the joint movement of varieties, size and input prices documented in Section 3.

4.4 A Simplified Case: Intensive and Extensive Margins

This subsection shows analytically how the search cost κ shapes the distribution of supplier networks, on the intensive margin of how many relationships an importer holds and on the extensive margin of whether a firm imports at all. The mechanism is easiest to see in a stripped down version of the model. I remove supplier heterogeneity by setting $z_u = \bar{z}_u$ for all potential suppliers, and set the bargaining weight to $\theta = 1$, which gives the buyer full bargaining power and thus a negotiated price equal to the seller's unit cost $1/\bar{z}_u$.

With no seller surplus, the contracting problem at a single match is

$$\max_{x \geq 0} \left[V(z, X_{\text{new}}) - V(z, X) - \tfrac{1}{\bar{z_u}}\, x \right],$$

where X is the current CES composite of inputs and $X_{\text{new}} = \left(X^\rho + x^\rho\right)^{1/\rho}$ after purchasing quantity x from the newly met supplier. For transparent comparative statics I solve a collapsed version that chooses the whole procurement path at once. This is equivalent to the sequential problem here, since supplier productivity is no longer random and the buyer holds all the bargaining power. Imposing symmetry across relationships, $x_i = x$ for all $i \in \{0, \dots, N\}$, the problem becomes

$$\max_{N \in \mathbb{N}_+,\, x \geq 0} \left\{ K_1 \left[z \left(\sum_{i=0}^{N} x_i^\rho \right)^{\frac{\alpha}{\rho}} \right]^{\frac{\epsilon - 1}{1 + \alpha(\epsilon - 1)}} - \tfrac{1}{\bar{z_u}}\, N x - N w \kappa \right\},$$

where K_1 collects constants and demand shifters, $\alpha \in (0, 1)$, and $\epsilon > 1$ is the demand elasticity in the final goods market. Symmetry implies $\sum_{i=0}^{N} x_i^\rho = N x^\rho$.

The optimal number of supplier searches satisfies

$$N^* = \left[\frac{K_2\, z^{\epsilon - 1}\, \bar{z_u}^{\alpha(\epsilon - 1)}}{w\, \kappa} \right]^{\phi}, \qquad \phi = \frac{\rho}{\rho - \alpha(\epsilon - 1)(1 - \rho)}.$$

where K_2 collects constants and demand shifters.[9] Throughout I assume $\rho - \alpha(\epsilon - 1)(1 - \rho) > 0$, so that $\phi > 0$.

N^* rises with buyer productivity z: more productive firms search more and hold larger networks. The elasticity of N^* with respect to the search cost is

$$\frac{\partial \ln N^*}{\partial \ln \kappa} = -\, \phi,$$

which is the same for every firm, so a proportional fall in κ raises all networks proportionally, and raises them most in levels for firms that already hold the most suppliers. Lower substitutability across inputs (lower ρ) makes each additional relationship more valuable and raises ϕ.

However, firms face an important constraint: N is an integer, so they adjust in steps

[9] $K_2 = \epsilon^{-\epsilon} \big(w / [(\epsilon - 1)(1 - \alpha)] \big)^{(1-\alpha)(1-\epsilon)} \big(C / P^{-\epsilon} \big) \big(\alpha(\epsilon - 1) \big)^{\alpha(\epsilon - 1)} \alpha(1 - \rho)/\rho$.

rather than smoothly. A firm holding many suppliers adds several at a time and tracks the continuous solution closely. A firm holding two or three only adjusts when it crosses the threshold at which one more relationship becomes profitable, so its network moves in jumps.

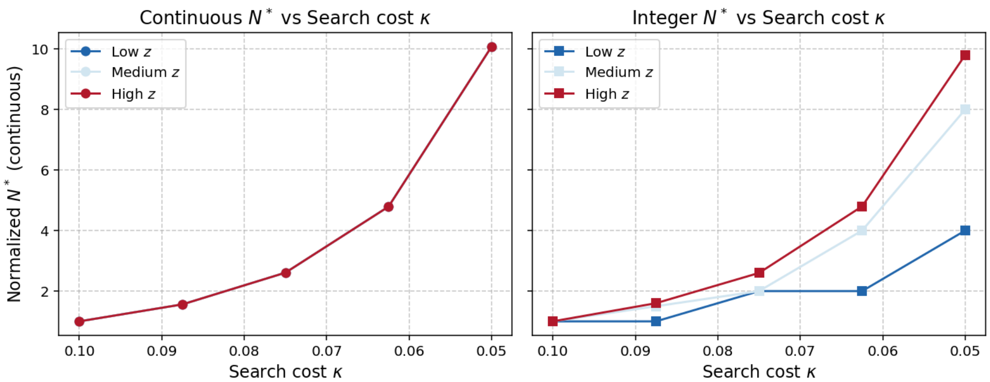

Figure 6: Discrete supplier choice generates lumpy adjustment for small firms and smoother adjustment for large firms.

A lower search cost also moves the extensive margin. Some small firms that previously chose zero imported relationships now find it profitable to enter.

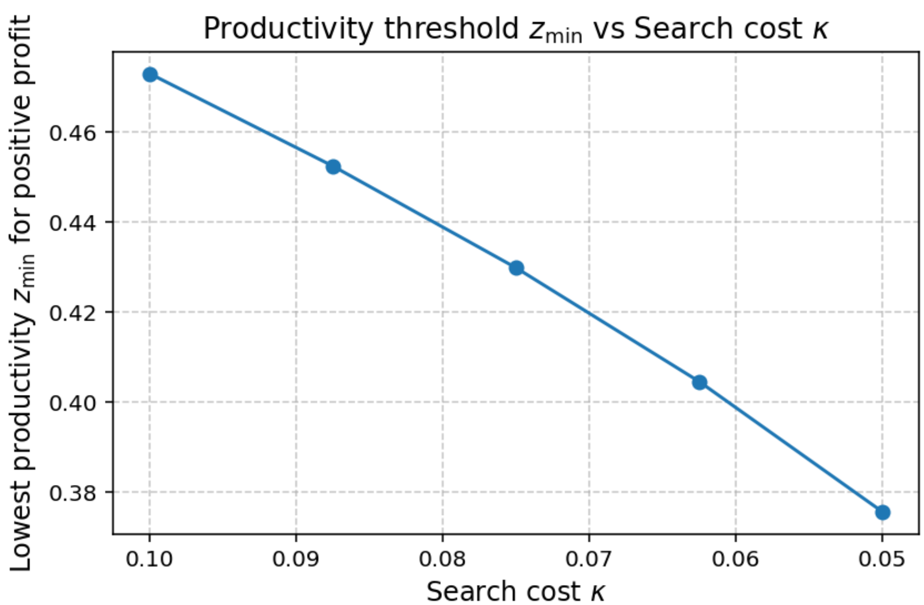

Figure 7: A decline in κ increases entry at the extensive margin.

A fall in κ therefore reshapes the distribution of supplier networks from both ends. On the intensive margin incumbents expand their networks, smoothly for large firms and in jumps for small ones. On the extensive margin new importers appear at the bottom. These are the two movements documented in Section 3.

Simplified Benchmark - Single Supplier Case works through a different simplification of the same model, one supplier drawn by random search from a heterogeneous pool, where the two margins reappear as a reservation supplier productivity and an entry threshold in z.

4.5 A Simplified Case: Price Dispersion

I now relax the assumption that the buyer has full bargaining power. Sellers receive weight $1-\theta$, and the input quantity is chosen efficiently to maximize the joint surplus. To isolate the role of bargaining, I continue to assume that all potential suppliers have identical productivity $\bar{z}_u$.

The efficient bargaining problem is

$$\max_{x\geq 0}\left(V(z,X_{\text{new}})-V(z,X)-\tfrac{1}{\bar{z}_u}x\right),$$

where $V(z,X_{\text{new}})$ is the buyer's continuation value after forming a new supplier relationship, and $V(z,X)$ is the buyer's outside option prior to the new match. The first-order condition defines the efficient input $x^*(z,X,\bar{z}_u)$.

Proposition 2. *Suppose $V_{Xz}(z,X_{\text{new}})>0$. With identical suppliers, the negotiated unit price is*

$$p_u=\frac{1-\theta}{\theta}\frac{w\kappa}{x^*}+\frac{1}{\bar{z}_u},$$

which decreases in buyer productivity z. Prices differ across buyers only through x^*, so the spread is proportional to $\frac{1-\theta}{\theta}w\kappa$ and vanishes when the buyer holds all the bargaining power or when search is costless.

Proof. Since $V_{Xz}(z,X_{\text{new}})>0$, the efficient input increases with buyer productivity for any fixed X:

$$\frac{\partial x^*(z,X,\bar{z}_u)}{\partial z}>0.$$

Every new search yields an identical supplier, so rejecting the current match and searching again leads to the same contract $(x^*(z,X,\bar{z}_u),T)$ at the cost of a further $w\kappa$. The buyer's outside option is therefore

$$V(z,X)=V(z,X_{\text{new}})-T-w\kappa.$$

Under Nash bargaining with weights $(\theta, 1-\theta)$ for the buyer and the seller, the transfer is

$$T = (1-\theta)\Big[V(z, X_{\text{new}}) - V(z, X)\Big] + \theta\,\frac{1}{\bar{z_u}}x^*.$$

Substituting the outside option gives $T = (1-\theta)(T + w\kappa) + \theta\,\frac{1}{\bar{z_u}}x^*$, or

$$T = \frac{1}{\bar{z_u}}x^* + \frac{1-\theta}{\theta}w\kappa.$$

Dividing through by x^* yields p_u. Since x^* increases in z, the first term falls in z while the second is common across buyers. $\square$

The payment splits into a variable component $\frac{1}{\bar{z_u}}x^*$, which covers the supplier's production cost, and a fixed component $\frac{1-\theta}{\theta}w\kappa$, which hands the seller a share of the buyer's surplus. Only the fixed component is spread over a larger quantity by more productive buyers, and that is what makes their unit price lower.

Both frictions are needed. With full buyer bargaining power the fixed component disappears, and with costless search there is nothing to bargain over beyond production cost. Since all suppliers here are identical, the dispersion the model generates cannot come from differences among them.

Bargaining is not the only source of the cost advantage. Because of love of variety in X, more productive firms hold more input varieties, so even with no dispersion in the unit price $\frac{\sum T}{\sum x}$ their effective price per unit of the composite, $\frac{\sum T}{X}$, is lower. CES aggregation therefore reinforces the advantage that bargaining creates.

See Sequential Search with Efficient Bargaining for further discussion of this sequential bargaining problem.

5 Quantitative Model and Results

This section takes the model to the data. I first add in-house production, so that firms have an alternative to importing, and close the model in general equilibrium. I then calibrate the decline in search costs to the observed growth in the number of imported varieties, and ask what that decline implies for the firm size distribution, the aggregate price level, and market concentration.

5.1 The Full Model: In-House Production

It would be unrealistic to assume that firms cannot produce at all without imported inputs. Each firm can instead produce an input in-house with a basic labor technology that does not depend on its own productivity, and because the domestic labor market is competitive, this input M is available at the unit labor cost $p_M = w$. After firms decide whether to search for foreign suppliers, production is

$$y = z \max\{X, M\}^{\alpha} l^{1-\alpha},$$

where X is the composite of imported inputs.

The max operator treats X and M as perfect substitutes, so a firm either sources from abroad or produces internally. A CES aggregator would let it do both, but at the cost of an additional continuous input-share choice, and it would not change the insight. The max specification keeps the model transparent and puts the weight on the discrete choice between importing and producing in-house.

5.2 General Equilibrium

The representative household supplies $\bar{L}$ units of labor inelastically and chooses consumption of each variety subject to its budget constraint,

$$\max_{c_i} \quad C = \left(\int c_i^{\frac{\epsilon-1}{\epsilon}} \, di \right)^{\frac{\epsilon}{\epsilon-1}} \quad \text{s.t.} \quad PC \leq w\bar{L} + \pi + \text{export},$$

where π is aggregate profits. The export term reflects that, since the economy imports some intermediate inputs, an equivalent value of domestic goods must be exported to balance trade. I assume these exports are sold at the domestic price P.

In equilibrium, aggregate labor demand equals the fixed labor supply,

$$\int \left(l_f + M_f + n_f \kappa \right) df = \bar{L},$$

where production labor l_f, intermediate production labor M_f and search effort $n_f \kappa$ depend on the demand shifter CP^{ϵ}.

The goods market clearing condition

$$\left(\int \mathbb{E}\left[y^{\frac{\epsilon-1}{\epsilon}} \,\middle|\, z \right] dz \right)^{\frac{\epsilon}{\epsilon-1}} = C$$

is automatically satisfied and therefore does not impose an additional equilibrium restriction. Aggregate demand C is by definition the CES aggregate of firm-level output, and each firm produces exactly the quantity demanded at its chosen price. Labor market clearing is therefore the only condition that has to be imposed, and it is what pins down CP^ϵ.

5.3 Numerical Solution and Simulation

To compute the equilibrium described above, I implement the nested fixed-point algorithm numerically. The outer loop, coded in Python, updates the aggregate demand term CP^ϵ until the labor market clears, while the inner loop, written in C++, solves the firm's dynamic problem and returns implied labor demand to the outer loop.

Inner Loop: Solving the Firm Problem. By Proposition 1, each firm is summarized by its productivity z and the composite X of its contracted inputs, and a new contract for x_{new} moves it to $X_{\text{new}} = (X^\rho + x_{\text{new}}^\rho)^{1/\rho}$. At each state (z, X), the firm decides whether to search for a new supplier. Conditional on meeting a supplier with productivity z_u, the two parties bargain over (x, T) to maximize the Nash product:

$$\max_{x,T} \left[V(z, X_{\text{new}}) - V(z, X) - T\right]^\theta \left[T - K(x)\right]^{1-\theta},$$

where $K(x) = x/z_u$ is the supplier's cost function.

Because this problem is strictly concave in the transfer T (see Concavity of the Bargaining Problem for a formal proof), the first-order condition with respect to T satisfies

$$(1-\theta)\left[V(z, X_{\text{new}}) - V(z, X) - T\right] = \theta\left[T - K(x)\right].$$

Solving for T gives the closed-form transfer rule

$$T = (1-\theta)V(x) + \theta K(x), \qquad V(x) = V(z, X_{\text{new}}) - V(z, X).$$

Substituting this expression back into the Nash product eliminates T and reduces the problem to a one-dimensional maximization:

$$\max_{x \geq 0} \theta^\theta (1-\theta)^{1-\theta} \Big[V(z, X_{\text{new}}) - V(z, X) - K(x)\Big].$$

This analytical reduction is crucial for tractability: it transforms a two-dimensional problem into a single-variable search over x. Because the objective is not globally concave in x due to the discrete search cost, I solve this problem by grid search rather than gradient-based optimization.

Given the optimal x^* and corresponding transfer T^*, the expected value of searching is updated as

$$V^S(z, X) = \mathbb{E}_{z_u}[V(z, X_{\text{new}}) - T^*(z, X, z_u)] - w\kappa,$$

and the overall value function is

$$V(z, X) = \max\{V^S(z, X), V^{NS}(z, X)\}.$$

Algorithmic Implementation. The inner loop is implemented in C++ and iterates on the value function over discretized grids for z and X. Because productivity enters the payoff multiplicatively and never enters the transition for X, the Bellman operator is block diagonal in z: each productivity type is a self-contained dynamic program. I therefore update one z row at a time, in place, and in parallel across rows.

One further feature is worth noting. Unlike a standard Bellman operator with discounting, the recursion here is affine, $V = \mathbb{E}[V'] - K$, with a constant search cost K. This mapping behaves like an identity transformation with a fixed shift, which converges monotonically but slowly. I therefore apply an adaptive over-relaxation step,

$$V^S_{t+1} = (1 - \omega_z)V^S_t + \omega_z \tilde{V}^S_t,$$

which "leans forward" toward the next iteration in the spirit of *successive over-relaxation (SOR)*, though applied here to a nonlinear value function mapping. Each row carries its own weight ω_z, raised while its residual keeps falling and cut back when progress stalls, which happens near the discrete boundary where firms switch between searching and not searching.

A row stops updating once its own sup-norm residual falls below a tight tolerance, and the inner loop ends when every row has converged. The resulting policy functions (x^*, T^*, V^S, V^{NS}) characterize the firm's optimal contracting and search behavior given aggregate conditions. Grid ranges and sizes, the treatment of the upper boundary, the convergence criteria, and the remaining implementation details are in Computational Details.

Simulation and Aggregation. After convergence, I simulate a large cross-section of firms to obtain the distribution of states. Each firm starts with $X = 0$ and draws productivity z from $F(z)$. In each simulated step, firms decide whether to search, draw suppliers z_u if searching, and form contracts according to (x^*, T^*). The input stock evolves as $X' = (X^\rho + x^{*\rho})^{1/\rho}$. The simulation produces firm-level histories of inputs, payments, and search activity, which are aggregated to compute total labor demand. The outer loop then updates CP^ϵ and repeats until the labor market clears.

Interpretation. The computational procedure combines a Python outer loop enforcing general equilibrium with a C++ inner loop solving firms' dynamic problems. The recursive state representation (z, X), the analytical elimination of T, and the adaptive over-relaxation scheme together make the model both tractable and faithful to the underlying theory.

5.4 Calibration

I take some parameters from external sources and calibrate the remainder by matching model moments to their data counterparts.

Table 2: Externally set parameters

Variable	Definition	Value(s)	Source
θ	Buyer bargaining power	0.83	Alviarez et al. (2023)
ϵ	Final market elasticity	4	Broda and Weinstein (2006)
ρ	Input elasticity	.75	Broda and Weinstein (2006)
α	Cobb–Douglas intermediate share	0.7	Intermediate share (manufacturing)
μ_{z_u}	Mean of supplier z_u	0	Normalization
σ_{z_u}	Std. of supplier z_u	0.235	Normalization

Table 3: Calibrated parameters and matched moments

Variable	Definition	Value(s)	Moment	Data	Model
μ_z	Mean of buyer z	0	Top 1% vs. median sales ratio	$259 \to 474$	$129 \to 211$
σ_z	Std. of buyer z	0.235	Top 1% sales share	$0.44 \to 0.49$	$0.424 \to 0.452$
κ_t	Search cost	$0.25 \to 0.15$	Mean number of varieties	$19.8 \to 24.4$	$19.7 \to 24.3$

Table 2 lists the parameters I take as given in the quantitative exercise. Buyer bargaining power is set to $\theta = 0.83$, taken from Alviarez et al. (2023). They develop a theory of market power that link buyers' market share, supplier market share and prices. Then, they use firm-to-firm trade data to estimate the bargaining power of U.S. firms. On the demand side, I adopt standard benchmark elasticities from Broda and Weinstein (2006): the final-goods market elasticity $\epsilon = 4$ and an input elasticity of substitution across varieties equal to 4. This is also consistent with the Swedish-specific elasticity from Boppart et al. (2023), where they find a median elasticity of 4.39 within industries. The Cobb–Douglas share on intermediates is $\alpha = 0.7$, chosen to match the intermediate-input expenditure share in manufacturing. For supplier productivity, I normalize the mean to the downstream scale, setting $\mu_{z_u} = 0$, and choose a dispersion $\sigma_{z_u} = 0.235$ to mirror the downstream distribution. None of these parameters are targeted in the calibration moments below; they serve as conventional benchmarks and normalizations.

Table 3 reports the calibration linking model primitives to sales and sourcing moments. To calibrate the moments, I begin with buyer productivity z, which I assume is lognormally distributed. This leaves two parameters to pin down, so I target two sales-based moments because physical productivity is difficult to measure accurately with available data. Specifically, within industries with at least five firms (to avoid mechanically concentrated sectors), I compute (i) the top-1% sales share and (ii) the ratio of top-1% sales to the median; I then take the median estimate across industries. Anchoring the location parameter at $\mu_z = 0$ gives a top-1% vs. median sales ratio that rises from $129 \rightarrow 211$ in the model, against $259 \rightarrow 474$ in the data. Setting the dispersion to $\sigma_z = 0.235$ matches the increase in the top-1% sales share ($0.424 \rightarrow 0.452$ in the model vs. $0.44 \rightarrow 0.49$ in the data). Given the lack of reliable quality- and price-adjusted productivity measures, using the sales distribution avoids additional structure (e.g., TFPR or labor-productivity corrections) and ties the two lognormal parameters to the level and the tail of the distribution, respectively. Finally, I discipline the time path of search frictions by lowering the per-period search cost from $\kappa_t : 0.25 \rightarrow 0.15$, which reproduces the observed growth in the mean number of imported varieties (data: $19.8 \rightarrow 24.4$; model: $19.7 \rightarrow 24.3$). This 40% decline is what the counterfactuals below impose: not an arbitrary shock, but the search-cost path implied by the variety data.[10]

5.5 Firm Size Dispersions

To study how search frictions affect the firm size distribution, I reduce the buyer search cost parameter by 40 percent in the calibrated model and recompute the stationary equilibrium. Lower search costs make it easier for firms to locate and bargain with additional suppliers. The resulting expansion in supplier networks is uneven across firms: more productive firms add more varieties, while some smaller firms begin importing.

[10]In the simulations I truncate the productivity distribution at the 0.01 and 99.99 percentiles.

The distributional consequence is a widening of firm size dispersion and, in the aggregate, higher market concentration.

Figure 8 shows the firm size distribution becoming more dispersed in both the data and the model: smaller firms begin to enter the market, while the most productive firms continue to expand rapidly. Section 4.4 gives the reason analytically. Adding one more supplier makes little difference to a firm that already has twenty, but going from one to two is a substantial jump for a smaller firm, so the most productive firms expand first and the fastest. As search costs decline further, less productive firms also begin to expand, but their growth eventually slows because the concavity of profits reduces the incentive to add more suppliers. However, bargaining and the love-of-variety structure of the intermediate-input aggregator counteract this force. Large firms with extensive supplier networks pay lower effective per-unit input costs, which reinforces their incentive to continue expanding their networks.

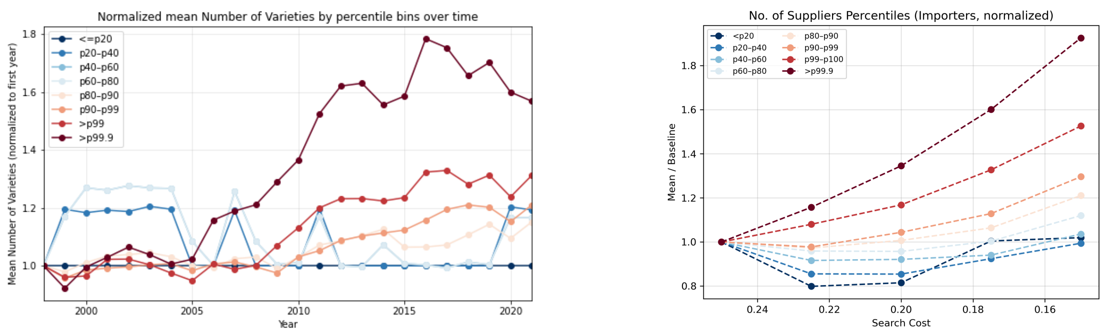

(a) Data (b) Model

Figure 8: Number of imported input varieties: data vs. model. Both show stronger supplier expansion among larger and more productive firms.

Both the data and the model show a right-tail expansion in the number of supplier relationships. This increase in supplier networks directly affects firm output and sales. Firms with more suppliers produce more efficiently and grow faster, widening the overall distribution of sales. Figure 9 illustrates this relationship.

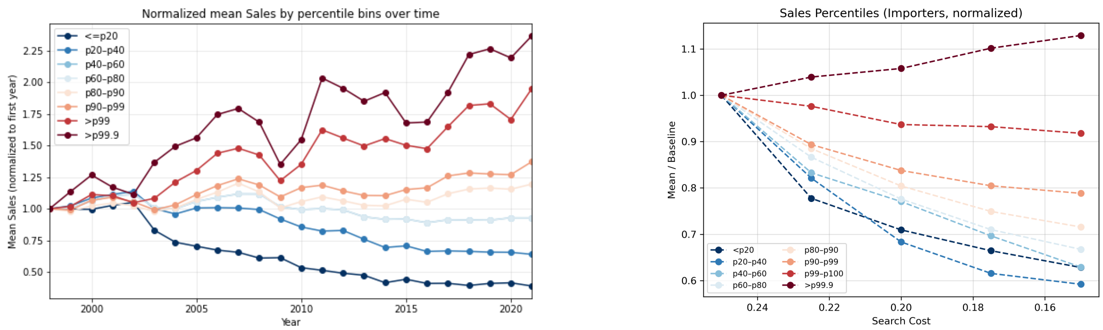

(a) Data (b) Model

Figure 9: Firm-size dispersion (sales): data vs. model. Both show widening dispersion over time, consistent with the mechanism in the model.

5.6 Price level, GDP and Concentration

Given my parameterization, the model quantifies how the decline in search costs affects aggregate outcomes: the price level, real GDP, and market concentration. I first define how market concentration is computed in real data and simulated data.

The model has a continuum of monopolistically competitive firms, so technically all firms are atomistic and concentration is zero by construction. A Herfindahl index is therefore not an object of the theory but a property of the simulated sample. I simulate 70,000 firms and compute the index over those with positive sales,

$$\text{HHI} = \sum_f s_f^2, \qquad s_f = \frac{\text{sales}_f}{\sum_g \text{sales}_g},$$

. Its level is governed by the number of simulated firms: with N equally sized firms, $\text{HHI} = 1/N$. Therefore, only proportional changes across counterfactuals are informative, and these are what I report.

In the data I treat a 5-digit industry as the market. For every 5-digit manufacturing sector I compute the HHI, and average across industries weighting by employment, so that the index describes the concentration facing the average worker. Since the model is assumed to be a single sector, this within-industry measure is the empirical counterpart.

The experiment implies a **13%** decline in the aggregate price level and a **13%** increase in real GDP. Labor supply is fixed, so this is also a 13% rise in output per worker, against roughly 49% growth in Sweden's real manufacturing value added per employee over the same period (SCB, 2025a,b,c,d). On the concentration side, the model delivers a

12.7% increase, against **27.8%** in the data between 1998 and 2021. Falling search costs therefore account for roughly **46%** of the observed rise in concentration, or between **32%** and **59%** if the comparison is instead made at the 2014 peak or after excluding the financial-crisis years.

Table 4: Aggregate effects when buyer search costs fall by **40 percent**

Price level and output	
Price level	−13%
Real GDP	+13%
Market concentration (HHI of sales)	
Data (overall, 1998–2021)	+27.8%
Data (peak, 2014)	+40.0%
Data (peak, excluding crisis years)	+21.5%
Model	+12.7%

Notes: Percent changes relative to the baseline calibrated equilibrium. Concentration is summarized by the Herfindahl–Hirschman index of firm sales, computed within 5-digit industries and averaged across them weighting by employment.

6 Testing the Mechanism: Broadband and Supplier Networks

The calibration in Section 5 disciplines the fall in search costs using a single aggregate moment, the mean number of imported varieties. This section brings independent variation to bear on the same mechanism. The staggered rollout of fiber broadband across Swedish municipalities changed the cost of finding and coordinating with suppliers at different times in different places, and none of that cross-sectional variation was used in the calibration.

The model makes two predictions here. Lower search costs raise the number of varieties a firm imports, and they raise it disproportionately for large firms, because a broader existing network makes the next supplier more valuable. The first is what the calibration was built to reproduce; the second is not, and it is the sharper test.

6.1 Fiber Coverage as a Proxy for Search Costs

I measure connectivity by the share of commercial establishments in a municipality with a fiber connection. Fiber is the fastest and most reliable broadband technology, supporting the video calls and real-time coordination that finding and managing a foreign supplier requires. Coverage is observed from 2010, when 27 percent of Swedish establishments had fiber, and reaches 79 percent by 2021. Its diffusion across municipalities is shown in Figure 29.

Let $\text{Varieties}_{f,t}$ denote the number of distinct imported input varieties used by firm f in year t. I estimate

$$\text{Varieties}_{f,t} = \beta_1 \,\text{FiberCoverage}_{m,t} + \beta_2 \left[\text{FiberCoverage}_{m,t} \times \text{BigFirm}_{f,t}\right] \\ + \text{FE}_t + \text{FE}_i + \text{FE}_m + \varepsilon_{f,t},$$

where $\text{FiberCoverage}_{m,t}$ is the fraction of establishments with fiber access in municipality m and year t, and $\text{BigFirm}_{f,t}$ indicates firms with sales above the contemporaneous industry mean, which is 18.8 percent of the sample. All regressions include year, industry, and municipality fixed effects, and standard errors are clustered by municipality. Firms with several establishments are assigned the connection of their best-connected one.

The sample consists of firms that import at least once, and retains the years in which such a firm imports nothing. Smaller importers buy intermittently, once every two or three years rather than annually, and dropping their zero years would discard exactly the year-to-year variation the design relies on. It would also select on the outcome.

Table 5: Connectivity and supplier-network size (levels specification)

	(1)	(2)
Fiber coverage	5.229***	0.307
	(1.004)	(0.200)
Fiber × BigFirm	—	23.314***
		(3.978)
BigFirm	—	17.882***
		(1.248)
Observations	349,400	349,400
R^2	0.024	0.070
Industry FE	Yes	Yes
Year FE	Yes	Yes
Municipality FE	Yes	Yes

Notes: Dependent variable is the level of imported input varieties for firm f in year t. "BigFirm" is an indicator equal to one for firms with sales above the contemporaneous industry mean. Standard errors are clustered by municipality. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Both predictions hold. Column 1 of Table 5 shows that firms in better-connected municipalities import more varieties. Column 2 shows where the effect sits: once the interaction is included, the coefficient on coverage alone is small and indistinguishable from zero, while the interaction with firm size is large and significant. Connectivity expands the networks of firms that already have large ones, which is the differential response the model generates and the calibration never targeted.

6.2 From Fiber Coverage to Search Costs

Comparing the estimates to the model requires putting the two on the same clock. I date the calibrated search-cost path by assuming κ falls linearly in time between 1998 and 2021, which is the assumption already implicit in comparing the model's five stationary equilibria to a time series. That places $\kappa = 0.198$ in 2010 and $\kappa = 0.150$ in 2021, a fall of 24 percent over the years for which fiber coverage is observed, and across that interval the model raises mean varieties among importers from 20.1 to 24.3.

The decline in κ and the diffusion of fiber are now changes over the same eleven years, but they are measured in different units. Varieties converts between them: the model says how many varieties its dated decline in κ delivers, the regression says how many the observed spread of fiber delivers, and the ratio of the two is the share of the model's expansion that broadband can account for.

The observed diffusion of fiber, from 27 to 79 percent coverage, implies $5.229 \times 0.52 \approx 2.7$ additional varieties. Over the same period the model predicts a gain of 4.2 varieties, so broadband accounts for 65 percent of that increase.

The large-firm margin gives a second and independent reading. For firms above their industry mean, the estimated effect is $(0.307 + 23.314) \times 0.52 \approx 12.3$ varieties, against 15.7 in the model for the top 18.8 percent of importers by sales: 79 percent of the model's expansion.

Two very different features of the estimates, the average effect and the gradient by size, therefore attribute 65 to 79 percent of the model's expansion to broadband alone. Neither had to agree with the other, and the model's figure exceeding the data's in both cases is the expected direction, since κ stands in for every source of falling search cost over these years and fiber is only one of them.

Two limitations bound how literally this should be read. The regression sample keeps the zero years of intermittent importers, while the model's mean is conditional on importing, because a firm that begins importing in the model never stops; supplier switching and inventory-driven intermittency are outside what the model is built to explain. And the coefficients are identified from differences across municipalities in a given year, whereas the calibrated decline in κ is economy-wide, so the exercise sets aside any general equilibrium response that the year effects absorb.

6.3 Robustness

Two further specifications are reported in the Appendix.

The first, in Robustness to Pretrend, replaces the continuous treatment with an indicator for municipal fiber coverage above 50 percent and adds a placebo dummy for the three years preceding the event. The concern it addresses is that municipalities with strong demand for connectivity, especially those hosting many large firms, might have received fiber earlier, in which case network expansion would reflect pre-existing local dynamics rather than connectivity itself. The main pre-event coefficient is close to zero and insignificant. The interaction $\textit{Preevent} \times \textit{BigFirm}$ is positive and significant at 5 percent, but at 3.49 it is under 15 percent of the corresponding post-treatment coefficient of 23.3. Large firms in better-connected municipalities may have been on a mild upward trend already, but not one large enough to account for the estimated effect.

The second, in Robustness Using Growth Rate Specification, re-estimates the relationship in growth rates rather than levels. The size gradient survives: the interaction remains positive and significant at 0.1817. This is worth stating plainly, because the simplified case in Section 4.4 implies a constant elasticity of network size to search costs, and therefore no size gradient once the outcome is normalized by firm size. That prediction fails. The full model does not impose constant elasticity, and the discrepancy is a reason to prefer it over the analytically convenient special case rather than a problem for the mechanism.

7 Rising Foreign Supplier Efficiency

The previous section examined how declining search costs improve input market efficiency and reallocate production toward more productive firms. A closely related secular trend has been the rise in foreign supplier productivity, especially among manufacturers in China and other major exporting economies. These two forces—lower search frictions and higher supplier efficiency—represent complementary aspects of the same transformation in global input markets.

This section quantifies the effects of rising supplier productivity and compares them directly with those of falling search costs under a calibration that yields nearly identical aggregate outcomes.

7.1 Calibration

To isolate the supplier-side mechanism, I vary only the mean productivity of foreign suppliers, denoted μ_{zu}. Supplier productivity starts from a relatively low level and rises until foreign suppliers reach the same average productivity as domestic firms. The search cost parameter is fixed at $\kappa = 0.15$, the level used in the comparison described in Section 5.

The magnitude of the productivity shift is chosen so that both experiments generate comparable changes in the model's demand shifter CP^{ϵ} and thus similar aggregate efficiency gains.

Table 6: Calibration for Rising Supplier Productivity Experiment

Parameter	Definition	Low Efficiency	High Efficiency
μ_{zu}	Mean supplier productivity	-0.15	0
κ	Search cost parameter	0.15	0.15

7.2 Aggregate Results

The rise in supplier productivity raises real GDP by about 10.5 percent and lowers the aggregate price level by roughly 11.4 percent. Market concentration, measured by HHI, increases by about 11.4 percent. These changes are very similar in magnitude to those generated by the fall in search costs considered earlier, indicating that both mechanisms produce comparable aggregate outcomes.

7.3 Firm-Level Outcomes

At the firm level, the simulated data show that the two mechanisms deliver similar changes in overall sales dispersion, but differ in the margin through which firms adjust.

When supplier efficiency rises, the number of input varieties falls on average and grows little for the most productive firms, consistent with efficiency improvements operating through existing relationships. The mean number of suppliers falls by about 22 percent, while the share of firms importing increases by 2.8 percentage points. Sales still rise broadly across the distribution. This decline in supplier count reflects substitution toward more efficient foreign partners rather than a reduction in openness, as firms rely more on a smaller set of high-productivity suppliers.

When search costs fall, by contrast, the number of input varieties expands substantially, reflecting the creation of new supplier matches. In short, the simulations indicate that rising supplier efficiency operates primarily through the intensive margin (greater trade per match), whereas lower search costs operate through the extensive margin (more matches).

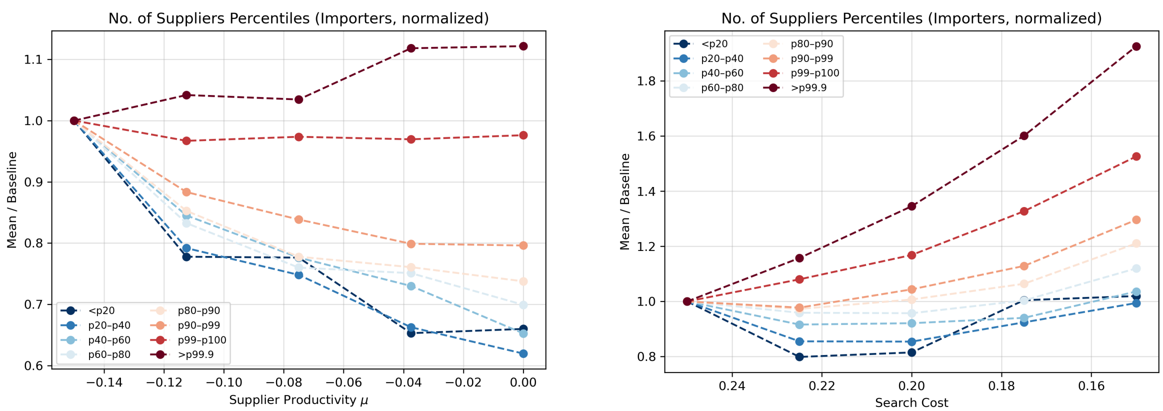

(a) Rising supplier efficiency

(b) Falling search costs

Figure 10: Number of input varieties across firm productivity percentiles. Rising supplier efficiency modestly changes variety counts, while falling search costs generate broad expansion through new matches.

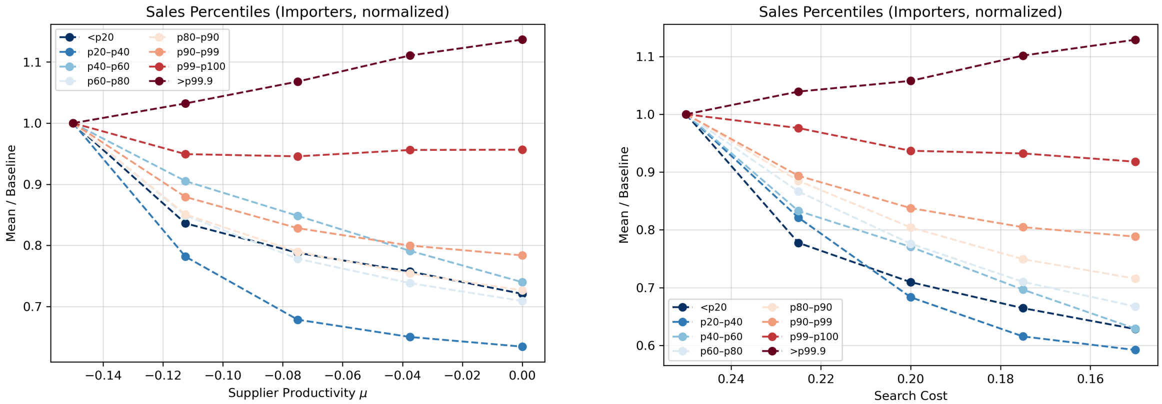

(a) Rising supplier efficiency

(b) Falling search costs

Figure 11: Firm sales across productivity percentiles. Both experiments raise sales across the distribution; the adjustment margin differs as described in the text.

7.4 Mechanism and Comparison

The simplified benchmark implies

$$N^* \propto \left(\frac{z\, z_u}{w\, \kappa}\right)^{\phi},$$

so both higher supplier productivity ($z_u\uparrow$) and lower search costs ($\kappa\downarrow$) raise the returns to forming input relationships.[11] Calibrating each shock to yield similar aggregate gains highlights that, in the data from these experiments, the principal difference lies in the margin of adjustment rather than in average efficiency.

Table 7: Comparison of Mechanisms at Matched Aggregate Gains

Outcome	Supplier Efficiency ↑	Search Cost ↓
Real GDP	+10.5%	+13%
Price level	−11.4%	−13%
Market concentration (HHI of sales)	+11.4%	+12.7%
Mean number of input varieties	−22%	+23%
Share of firms importing	+2.8 pp	+3.3 pp

Overall, both mechanisms deliver similar aggregate outcomes and similar sales dispersion. The distinguishing pattern in these simulations is the composition of adjustment: lower search costs expand the set of input varieties used by firms, while higher supplier efficiency raises the intensity of trade within a smaller set of varieties.

[11]This expression is a shorthand for the full benchmark derived in Section 4.4.

8 Counterfactual: Tariff

Another dimension of change in input markets comes from trade policy and institutional barriers to trade. Tariffs, export controls, and localization requirements have fluctuated over time, but recent years have seen a renewed rise in trade restrictions as part of a broader policy turn toward deglobalization or fragmentation (e.g. WTO, 2023; Aiyar et al., 2023). These developments directly affect firms' access to foreign suppliers and the relative cost of imported intermediates. In this section, I examine how such policy changes, in particular a tariff on imported inputs, alter key outcomes including productivity, market concentration, and aggregate welfare. The analysis connects technological changes in input market efficiency to tangible policy instruments that affect firms in similar ways.

8.1 Policy and Implementation

I introduce an ad valorem tariff on imported intermediate inputs at rate $\tau = 0.10$. Tariff revenue is rebated to the representative household as a lump-sum transfer. The tariff scales the buyer's payment to foreign suppliers by the factor $(1+\tau)$, so that the buyer's effective unit input price becomes

$$p_u^{\text{imp}} = (1+\tau)\, p_u.$$

Given this wedge, the Nash bargaining problem between buyer and supplier becomes

$$\max_{T,\, x} \Big(-(1+\tau)T + V(z, X_{\text{new}}) - V(z, X) \Big)^{\theta} \Big(T - w_f \tfrac{x}{z_u} \Big)^{1-\theta}.$$

The tariff therefore raises the cost of imported inputs proportionally to $(1+\tau)$ while leaving the rest of the environment, including the production structure and demand system, unchanged.

8.2 Aggregate Effects

The tariff increases the effective cost of imported inputs, which feeds through the CES aggregator into higher consumer prices. Figure 12 compares aggregate outcomes under the baseline and a 10% tariff across different search costs. In the baseline, lower search costs reduce prices and raise production, the domestic input share, and real GDP. With the tariff, these relationships remain qualitatively similar but are uniformly dampened: prices are higher, production and real GDP are lower, and the shift toward domestic inputs is stronger. The tariff thus partially reverses the aggregate gains from improved supplier search efficiency.

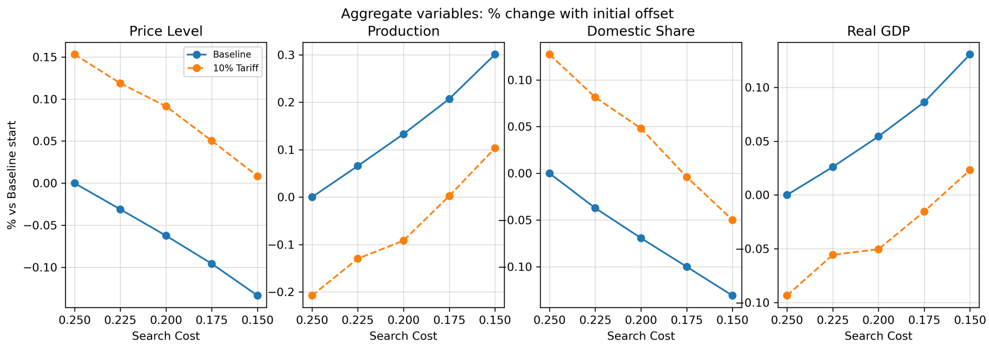

Figure 12: Aggregate outcomes under a 10% tariff. The tariff raises prices and weakens the production and real GDP response to lower search costs, while increasing the domestic input share.

8.3 Search Incentives and Network Size

By scaling the variable cost of imported inputs, the tariff reduces the gain from adding a supplier. In the simplified benchmark with $z_u = \bar{z}_u$ and buyer power $\theta = 1$, a tariff τ scales the supplier's unit cost by $(1+\tau)$, equivalent to replacing $\bar{z}_u$ with $\bar{z}_u/(1+\tau)$. Since $N^* \propto \bar{z}_u^{\alpha(\epsilon-1)\phi}$ in Section 4.4, the number of supplier relationships under a tariff becomes

$$N_\tau^* = N_0^*(1+\tau)^{-\alpha(\epsilon-1)\phi}, \qquad \phi = \frac{\rho}{\rho - \alpha(\epsilon-1)(1-\rho)}, \tag{2}$$

where N_0^* denotes the no-tariff benchmark. The model therefore predicts a broad decline in network size when the tariff is introduced. This effect is confirmed in the full quantitative model, as shown below.

Table 8: Change in number of suppliers, baseline versus tariff, both at low search cost.

	Share (%)	**Baseline**	**Tariff**
Fewer	95.87	25.08	17.70
More	0.00	*n.a.*	*n.a.*
Equal	4.13	6.16	6.16

Notes: Firms are grouped by whether the tariff leaves them with fewer, more, or the same number of suppliers. "Share" is the percentage of firms in each group; the remaining columns give that group's mean number of suppliers with and without the tariff. Both scenarios are evaluated at the low search cost, $\kappa = 0.15$.

8.4 Distributional Effects on Sales and Profits

The tariff changes relative prices and demand. Nearly all firms end up with higher sales and profits, benefiting from higher output prices, while the small minority that loses is made up of much larger firms, whose imported inputs have become more expensive. I report the share of firms with lower or higher outcomes and the corresponding averages.

Table 9: Sales and profits under the baseline and the tariff (low search cost).

Panel A: Sales

	Share (%)	Baseline	Tariff
Lower	5.58	27.81	25.45
Higher	94.42	0.69	0.83

Panel B: Profits

	Share (%)	Baseline	Tariff
Lower	3.73	2.08	1.84
Higher	96.27	0.16	0.20

Notes: Firms are grouped by whether the tariff leaves them with lower or higher sales (Panel A) or profits (Panel B). "Share" is the percentage of firms in each group; the remaining columns give that group's mean with and without the tariff. Both scenarios are evaluated at the low search cost, $\kappa = 0.15$.

8.5 Market Concentration

I measure concentration using the HHI, computed over the simulated sample as described in Section 5, and the sales gap between the largest and smallest importing firms. The tariff compresses the firm size distribution in the model, both in general and among importers. Larger firms, which import the most, have less incentive to sustain a wide supplier network once imported inputs become more expensive.

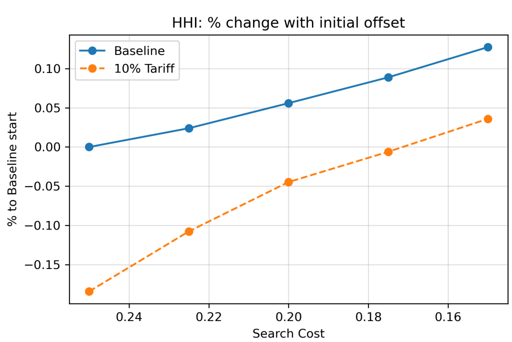

(a) HHI of sales.

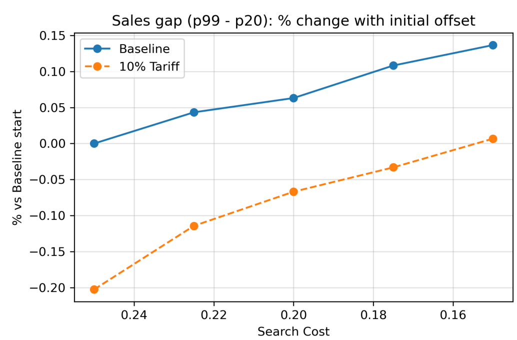

(b) Sales gap between top and bottom importing firms.

Figure 13: Market concentration and dispersion of sales under the baseline and tariff scenarios.

To make the comparison explicit, I contrast the low-search-cost equilibrium with a 10% tariff to the high-search-cost baseline. Relative to the baseline ($\kappa = 0.25$), the low-search-cost case ($\kappa = 0.15$) with a 10% tariff delivers real GDP higher by 6.98% and an HHI higher by 3.60%. For reference, moving from high to low search cost with no tariff raises real GDP by 13.06% and HHI by 12.74%. Adding the 10% tariff at low search cost therefore reduces GDP by 5.37% and HHI by 8.11% relative to that low-search-cost, no-tariff case.

Overall, the tariff counterfactual shows that protectionist policies can mimic several features of higher search frictions by discouraging supplier formation and compressing firm networks. Although such policies reduce aggregate productivity, they may appear politically attractive because they narrow the performance gap between large and small firms and limit further spatial concentration of economic activity.

9 Conclusion

This paper studies how frictions in input markets shape the firm size distribution. I document three dimensions along which Swedish importing firms differ from one another, and increasingly so: the number of varieties they import, their sales, and the prices they pay for inputs. Motivated by these facts, I build a quantitative model of supplier search and bargaining. A 40% fall in search costs accounts for 27% of measured labor productivity growth and just under half of the rise in market concentration since 1998, and the mechanism is validated against the staggered rollout of high-speed internet.

The framework brings together features that are usually studied apart: buyers with networks of very different sizes, suppliers of differing efficiency, and prices negotiated one relationship at a time. Having a unified framework is important since any change in the cost of reaching a supplier affects these margins, and I show that the interplay of these margins can help explain important empirical patterns, such as the dispersion of input prices. This framework also matters for many policies, not only the tariff shown in Section 8, since any instrument that affects different margins of the input market can be accommodated easily into this framework. More importantly, the quantification of this model shows that input market frictions are a first-order force for macroeconomic outcomes.

A further contribution is methodological. I show how the CES structure, which is common for production functions, demand systems and utility functions, can be used easily to collapse a portfolio of arbitrarily many heterogeneous suppliers into a single composite in dynamic problems. In the case of this paper, the state stays two-dimensional however large a network becomes. This is a result that can be applied to many fields of structural economics: any dynamic problem in which an agent accumulates a variable number of heterogeneous objects, seen only through a CES aggregate, admits the same reduction at the cost of one line of algebra.

This paper has important implications for many economic problems beyond the immediate trade and input market context, such as spatial ones. For example, the mechanisms identified here may also contribute to the spatial concentration of economic activity, since larger and more productive firms tend to cluster in major urban regions across the EU, where such patterns have intensified over time (OECD, 2023; European Commission Directorate General for Regional and Urban Policy, 2024). Policies that influence search costs, such as infrastructure investment, digital connectivity, or trade regulation, can materially affect spatial market structure. Extending the model to study these policies and patterns can be an interesting direction for future research.

References

Anuraag Aekka, NYU Stern, and Gaurav Khanna. Endogenous Production Networks and Firm Dynamics. 2026.

Philippe Aghion, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li. A Theory of Falling Growth and Rising Rents. *The Review of Economic Studies*, 90(6): 2675–2702, November 2023. ISSN 0034-6527. doi: 10.1093/restud/rdad016. URL https://doi.org/10.1093/restud/rdad016.

Shekhar Aiyar, Anna Ilyina, Jiaqian Chen, Alvar Kangur, Juan Trevino, Christian Ebeke, Tryggvi Gudmundsson, Gabriel Soderberg, Tatjana Schulze, Tansaya Kunaratskul, Michele Ruta, Roberto Garcia-Saltos, and Sergio Rodriguez. Geo-Economic Fragmentation and the Future of Multilateralism. *Staff Discussion Notes*, 2023 (001):1, January 2023. ISSN 2617-6750. doi: 10.5089/9798400229046.006. URL http://dx.doi.org/10.5089/9798400229046.006.

George Alessandria, Joseph P. Kaboski, and Virgiliu Midrigan. Inventories, Lumpy Trade, and Large Devaluations. *American Economic Review*, 100(5):2304–2339, December 2010. ISSN 0002-8282. doi: 10.1257/aer.100.5.2304. URL https://www.aeaweb.org/articles?id=10.1257/aer.100.5.2304.

Vanessa I. Alviarez, Michele Fioretti, Ken Kikkawa, and Monica Morlacco. Two-Sided Market Power in Firm-to-Firm Trade, May 2023. URL https://www.nber.org/papers/w31253.

Mary Amiti and Sebastian Heise. U.S. Market Concentration and Import Competition. *The Review of Economic Studies*, 92(2):737–771, March 2025. ISSN 0034-6527. doi: 10.1093/restud/rdae045. URL https://doi.org/10.1093/restud/rdae045.

Mary Amiti and Jozef Konings. Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *American Economic Review*, 97(5):1611–1638, December 2007. ISSN 0002-8282. doi: 10.1257/aer.97.5.1611. URL https://www.aeaweb.org/articles?id=10.1257/aer.97.5.1611.

Pol Antràs, Teresa C. Fort, and Felix Tintelnot. The Margins of Global Sourcing: Theory and Evidence from US Firms. *American Economic Review*, 107(9):2514–2564, September 2017. ISSN 0002-8282. doi: 10.1257/aer.20141685. URL https://www.aeaweb.org/articles?id=10.1257/aer.20141685.

Enghin Atalay. Materials Prices and Productivity. *Journal of the European Economic Association*, 12(3):575–611, 2014. ISSN 1542-4766. URL https://www.jstor.org/stable/90023276.

Matej Bajgar, Chiara Criscuolo, and Jonathan Timmis. Intangibles and Industry Concentration: A Cross-Country Analysis. *Oxford Bulletin of Economics and Statistics*, n/a(n/a), 2025. ISSN 1468-0084. doi: 10.1111/obes.12659. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/obes.12659. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/obes.12659.

Johannes Boehm and Ezra Oberfield. Misallocation in the Market for Inputs: Enforcement and the Organization of Production*. *The Quarterly Journal of Economics*, 135(4):2007–2058, November 2020. ISSN 0033-5533. doi: 10.1093/qje/qjaa020. URL https://doi.org/10.1093/qje/qjaa020.

Timo Boppart, Mikael Carlsson, Markus Kondziella, and Markus Peters. Micro PPI-Based Real Output Forensics. 2023.

Christian Broda and David E. Weinstein. Globalization and the Gains From Variety*. *The Quarterly Journal of Economics*, 121(2):541–585, May 2006. ISSN 0033-5533. doi: 10.1162/qjec.2006.121.2.541. URL https://doi.org/10.1162/qjec.2006.121.2.541.

Ariel Burstein, Javier Cravino, and Marco Rojas. Input Price Dispersion Across Buyers and Misallocation, November 2024. URL https://www.nber.org/papers/w33128.

Matias Covarrubias, Germán Gutiérrez, and Thomas Philippon. From Good to Bad Concentration? US Industries over the Past 30 Years. *NBER Macroeconomics*

Annual, 34:1–46, January 2020. ISSN 0889-3365. doi: 10.1086/707169. URL https://www.journals.uchicago.edu/doi/full/10.1086/707169.

Nicolas Crouzet and Janice C. Eberly. Understanding Weak Capital Investment: the Role of Market Concentration and Intangibles. Technical Report w25869, National Bureau of Economic Research, May 2019. URL https://www.nber.org/papers/w25869.

Maarten De Ridder. Market Power and Innovation in the Intangible Economy. *American Economic Review*, 114(1):199–251, January 2024. ISSN 0002-8282. doi: 10.1257/aer.20201079. URL https://www.aeaweb.org/articles?id=10.1257/aer.20201079.

Angus Deaton. Quality, Quantity, and Spatial Variation of Price. *The American Economic Review*, 78(3):418–430, 1988. ISSN 0002-8282. URL https://www.jstor.org/stable/1809142.

Swati Dhingra and John Morrow. Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity. *Journal of Political Economy*, 127(1):196–232, February 2019. ISSN 0022-3808. doi: 10.1086/700732. URL https://www.journals.uchicago.edu/doi/full/10.1086/700732.

Jonathan Eaton, Samuel Kortum, and Francis Kramarz. An Anatomy of International Trade: Evidence From French Firms. *Econometrica*, 79(5):1453–1498, 2011. ISSN 1468-0262. doi: 10.3982/ECTA8318. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA8318. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA8318.

Jonathan Eaton, David Jinkins, James R. Tybout, and Daniel Xu. Two-Sided Search in International Markets, January 2022a. URL https://www.nber.org/papers/w29684.

Jonathan Eaton, Samuel S Kortum, and Francis Kramarz. Firm-to-Firm Trade: Imports, Exports, and the Labor Market. 2022b.

Chris Edmond, Virgiliu Midrigan, and Daniel Yi Xu. Competition, Markups, and the Gains from International Trade. *American Economic Review*, 105(10):3183–3221, October 2015. ISSN 0002-8282. doi: 10.1257/aer.20120549. URL https://www.aeaweb.org/articles?id=10.1257/aer.20120549.

Jan Eeckhout and Philipp Kircher. Sorting and Decentralized Price Competition. *Econometrica*, 78(2):539–574, 2010. ISSN 1468-0262. doi: 10.3982/ECTA7953. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA7953. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA7953.

Wilfred J. Ethier. National and International Returns to Scale in the Modern Theory of International Trade. *The American Economic Review*, 72(3):389–405, 1982. ISSN 0002-8282. URL https://www.jstor.org/stable/1831539.

European Commission Directorate General for Regional and Urban Policy. *Ninth report on economic, social and territorial cohesion*. Publications Office, LU, 2024. URL https://data.europa.eu/doi/10.2776/264833.

Takeshi Fukasawa. The use of symmetry for models with variable-size variables. *Mathematical Social Sciences*, 138:102437, December 2025. ISSN 0165-4896. doi: 10.1016/j.mathsocsci.2025.102437. URL https://www.sciencedirect.com/science/article/pii/S0165489625000526.

Pinelopi Koujianou Goldberg, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova. Imported Intermediate Inputs and Domestic Product Growth: Evidence from India*. *The Quarterly Journal of Economics*, 125(4):1727–1767, November 2010. ISSN 0033-5533. doi: 10.1162/qjec.2010.125.4.1727. URL https://doi.org/10.1162/qjec.2010.125.4.1727.

Gita Gopinath and Brent Neiman. Trade Adjustment and Productivity in Large Crises. *American Economic Review*, 104(3):793–831, March 2014. ISSN 0002-8282. doi: 10.1257/aer.104.3.793. URL https://www.aeaweb.org/articles?id=10.1257/aer.104.3.793.

Gautam Gowrisankaran and Marc Rysman. Dynamics of Consumer Demand for New Durable Goods. *Journal of Political Economy*, 120(6):1173–1219, December 2012. doi: 10.1086/669540. URL https://www.journals.uchicago.edu/doi/abs/10.1086/669540.

Gene M. Grossman, Elhanan Helpman, and Alejandro Sabal. Resilience in Vertical Supply Chains, September 2023. URL https://www.nber.org/papers/w31739.

Gene M. Grossman, Elhanan Helpman, and Stephen J. Redding. When Tariffs Disrupt Global Supply Chains. *American Economic Review*, 114(4):988–1029, April 2024. ISSN 0002-8282. doi: 10.1257/aer.20211519. URL https://www.aeaweb.org/articles?id=10.1257/aer.20211519.

László Halpern, Miklós Koren, and Adam Szeidl. Imported Inputs and Productivity. *American Economic Review*, 105(12):3660–3703, December 2015. ISSN 0002-8282. doi: 10.1257/aer.20150443. URL https://www.aeaweb.org/articles?id=10.1257%2Faer.20150443&utm_source=chatgpt.com.

Chang-Tai Hsieh and Peter J. Klenow. Misallocation and Manufacturing TFP in China and India*. *The Quarterly Journal of Economics*, 124(4):1403–1448, November 2009. ISSN 0033-5533. doi: 10.1162/qjec.2009.124.4.1403. URL https://doi.org/10.1162/qjec.2009.124.4.1403.

Maurice Kugler and Eric Verhoogen. Prices, Plant Size, and Product Quality. *The Review of Economic Studies*, 79(1):307–339, January 2012. ISSN 0034-6527. doi: 10.1093/restud/rdr021. URL https://doi.org/10.1093/restud/rdr021.

Jeremy Lise, Costas Meghir, and Jean-Marc Robin. Matching, sorting and wages. *Review of Economic Dynamics*, 19:63–87, January 2016. ISSN 1094-2025. doi: 10.1016/j.red.2015.11.004. URL https://www.sciencedirect.com/science/article/pii/S109420251500071X.

Julien Martin, Isabelle Mejean, and Mathieu Parenti. Relationship stickiness and economic uncertainty. 2023. URL http://www.isabellemejean.com/MMP_2023.pdf.

J. J. McCall. Economics of Information and Job Search. *The Quarterly Journal of Economics*, 84(1):113–126, 1970. ISSN 0033-5533. doi: 10.2307/1879403. URL https://www.jstor.org/stable/1879403.

Sébastien Miroudot, Rainer Lanz, and Ragoussis Ragoussis. Trade in Intermediate Goods and Services, November 2009. URL https://www.oecd.org/en/publications/trade-in-intermediate-goods-and-services_5kmlcxtdlk8r-en.html.

Monica Morlacco. Market Power in Input Markets: Theory and Evidence from French Manufacturing, 2020. URL https://www.dropbox.com/s/4ic4ylvgngzn947/BP_draft_01102020.pdf?dl=0.

Ezra Oberfield. A Theory of Input-Output Architecture. *Econometrica*, 86(2):559–589, 2018. ISSN 1468-0262. doi: 10.3982/ECTA10731. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA10731. _eprint: https://onlinelibrary.wiley.com/-doi/pdf/10.3982/ECTA10731.

OECD. *OECD Regional Outlook 2023: The Longstanding Geography of Inequalities*. OECD Regional Outlook. OECD Publishing, October 2023. ISBN 978-92-64-81264-2 978-92-64-32839-6 978-92-64-34116-6. doi: 10.1787/92cd40a0-en. URL https://www.oecd.org/en/publications/oecd-regional-outlook-2023_92cd40a0-en.html.

Ministry of Enterprise and Innovation. A Completely Connected Sweden by 2025, a Broadband Strategy. Technical Report N2016/08008/D, Ministry of Enterprise and Innovation, Sweden, December 2016.

Christopher A. Pissarides. *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, USA, 2 edition, February 2017. ISBN 978-0-262-53398-0.

James E. Rauch. Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35, June 1999. ISSN 0022-1996. doi: 10.1016/S0022-1996(98)00009-9. URL https://www.sciencedirect.com/science/article/pii/S0022199698000099.

Diego Restuccia and Richard Rogerson. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720, October 2008. ISSN 1094-2025. doi: 10.1016/j.red.2008.05.002. URL https://www.sciencedirect.com/science/article/pii/S1094202508000203.

Michael Rubens. Market Structure, Oligopsony Power, and Productivity. *American Economic Review*, 113(9):2382–2410, September 2023. ISSN 0002-8282. doi: 10.1257/aer.20210383. URL https://www.aeaweb.org/articles?id=10.1257/aer.20210383.

SCB. STATISTIKENS FRAMSTAELLNING, Utrikeshandel med varor, 2018. February 2018.

SCB. Antal sysselsatta (ENS2010), kalenderkorrigerade och säsongrensade, 1000-tal efter näringsgren SNI 2007 och kvartal. PxWeb, 2025a. URL https://www.statistikdatabasen.scb.se/pxweb/sv/ssd/START__NR__NR0103__NR0103B/NR0103ENS2010T22KvN/table/tableViewLayout1/.

SCB. Industrial production index, not calendar adjusted, chain index, 2000=100 by industrial classification NACE Rev. 1.1 and year. PxWeb, 2025b. URL https://www.statistikdatabasen.scb.se/pxweb/en/ssd/START__NV__NV0402__NV0402C/IPI2000kedjAr/table/tableViewLayout1/.

SCB. Industrial production index, not calendar adjusted, chain index, 2021=100 by industrial classification (NACE Rev. 2) and month. PxWeb, 2025c. URL https://www.statistikdatabasen.scb.se/pxweb/en/ssd/START__NV__NV0402__NV0402A/IPI2010KedjM/table/tableViewLayout1/.

SCB. Output, intermediate consumption and value added (ESA2010), current prices, SEK million by industrial classification (NACE Rev. 2), transaction item and year. PxWeb, 2025d. URL https://www.statistikdatabasen.scb.se/pxweb/en/ssd/START__NR__NR0103__NR0103E/NR0103ENS2010T09RK/table/tableViewLayout1/.

Robert Shimer and Lones Smith. Assortative Matching and Search. *Econometrica*, 68(2): 343–369, 2000. ISSN 0012-9682. URL https://www.jstor.org/stable/2999430.

Joshua Weiss. Intangible Investment and Market Concentration. 2020. URL https://jweissecon.github.io/JobMarket/WeissPaper.pdf.

WTO. World Trade Report 2023 - Re-globalization for a secure, inclusive and sustainable future. Technical report, World Trade Organization, 2023. URL https://www.wto.org/english/res_e/publications_e/wtr23_e.htm.

A Data Appendix

A.1 Summary Statistic

The following summary statistics are computed for an average year in my dataset:

Table 10: Summary Statistics — Imports

Variable	Mean	Std. Dev.	P10	Median	P95
No. Varieties per Firm	20.8	78.7	1.0	4.0	84.7
No. Product per Firm	13.2	31.6	1.0	3.3	54.8
No. Country per Firm	4.6	6.3	1.0	2.0	17.0
No. Country per Product	7.7	8.6	1.0	5.0	24.9
No. Firms per Variety	3.1	7.3	1.0	1.0	10.0
No. Country per FirmXProduct	1.6	1.8	1.0	1.0	4.0
No. of Firms			8396		
No. of Product			7400		
No. of Varieties			56608		
No. of Countries			179		
Percentage of Domestic Firms			0.28		
Imported Intermediate share			0.02383		

The next table compares importing firms with all firms:

Table 11: Summary Statistics — Importing Firm vs All Firm

Variable	Importers	All Firms	Ratio
No. of Firms per Year	8396	52811	0.16
Total No. of Firms	36483	150046	0.24
Sales	1.63e+08	2.83e+07	5.77
Employment	59.06	11.15	5.30
Labor Productivity	13.72	12.82	1.07
Intermediate share of Production	0.44	0.36	1.23

A.1.1 Within EU Data Selection and Non-EU subsample

Imports originating from other EU countries do not pass through customs, since Sweden has been part of the EU since 1995. Data on EU imports are therefore collected from mandatory self-reports by firms. A firm importing below a certain value threshold has no obligation to report, however. This cutoff is 1.5 million SEK worth of goods in 1998–2004, 2.2 million SEK in 2005–2008, 4.5 million SEK in 2009–2014 and 9.0 million

SEK thereafter (see SCB (2018) for more). Because of these rules, Britain's exit from the EU in 2020 also causes some irregularities in the data for 2021. Limiting the sample to non-EU imports, or excluding Britain in all years, does not change the conclusions of this paper by much. Focusing on the non-EU sample also gives a better view of the bottom of the distribution of importing firms. Non-EU imports account for around 50% of the full sample.

I include robustness tests using only non-EU imports for several of my empirical facts. I define non-EU imports as imports from countries that were never in the EU throughout the sample period. Therefore, for example, import records from Bulgaria (joined EU in 2007) in 2001 will be excluded from those robustness tests.

A.1.2 Productivity Measurement

I first compute revenue-based productivity (TFPR) under a Cobb–Douglas production function:

$$\mathrm{TFPR}_{f,t} = \frac{p_{f,t} y_{f,t}}{\left(w_{f,t}\,\ell_{f,t}\right)^{\alpha_\ell} K_{f,t}^{\alpha_k} \left(\sum_u T_{f,t,u}\right)^{1-\alpha_\ell-\alpha_k}}.$$

Here, $p_{f,t} y_{f,t}$ denotes firm f's nominal output (price times quantity) in year t, $w_{f,t}\ell_{f,t}$ is the wage bill, $K_{f,t}$ is the capital stock, and $T_{f,t,u}$ are intermediate input expenditures (indexed by input type u). All variables are observed in firms' financial statements.

I set the rental rate of capital to $r = 0.15$, following standard values in the literature. The elasticities α_ℓ and α_k are obtained using the Cobb–Douglas cost-share property: under cost minimization and competitive input markets, each elasticity equals the input's expenditure share. I assume these elasticities are constant within an industry and equal to the industry's average expenditure shares. Provided the firm uses strictly positive labor, intermediate inputs, and capital, the TFPR measure is well defined.

I then estimate physical productivity (TFPQ) as:

$$\mathrm{TFPQ}_{f,t} = \frac{\mathrm{TFPR}_{f,t}}{p_{f,t}}.$$

Here, $p_{f,t}$ denotes the firm-level output price. I obtain $p_{f,t}$ from the producer price index (PPI) or from an export database. A limitation is that PPI coverage is concentrated among large firms, which are often exporters, so this sample is biased toward larger firms. In Quality-adjusted TFPQ, I construct model-based output prices that adjust for product quality and allow for a broader set of firms to be included. The conclusions are robust to using these model-based prices.

A.1.3 Robustness Check — Arm length

To test whether my results hold for arm-length trade, that is, trade that does not take place within the same multinational corporate group, I carry out robustness tests on the sub-sample of firms that neither (1) belong to a foreign-based corporate group nor (2) belong to a corporate group that owns subsidiaries abroad. Any international trade by these firms must therefore be at arm's length. This subsample contains around 25% of the observations in the original dataset.

For arm-length trade, (1) the negative relationship between TFPQ and input prices is unchanged; (2) the evolution of the number of imported varieties is slightly different, with the 99th percentile firm still increasing significantly and the 90th to 99th percentiles dropping slightly while the rest are flat, though I believe the same conclusion can be drawn; (3) the number of imported varieties rises much less with TFPQ once these companies are excluded, which is not surprising, and it is unclear whether this reflects the exclusion of non-arm-length trade or of the largest and most productive multinationals; (4) the price dispersion result also holds, with the absolute value of the variance increasing in all years. Interestingly, there is a large drop in 2020 (COVID-19) that is not seen in the original graph.

A.2 Robustness Tests on Number of Variety Percentiles Evolution

The main text ranks firms by the number of varieties they import. Here I instead measure firm size by annual sales, and collect further robustness checks on how the variety distribution has evolved. A variety remains a unique country-product pair, and each percentile is normalized by its 1998 value. While the median and lower-percentile firms have experienced no change or even a slight decrease in the number of imported varieties from 1998 to 2021, the largest firms have increased their imported varieties by approximately 10% to 30%. Since the number of varieties a firm imports reflects its search activity, this points to differential effects of search costs across firm sizes.

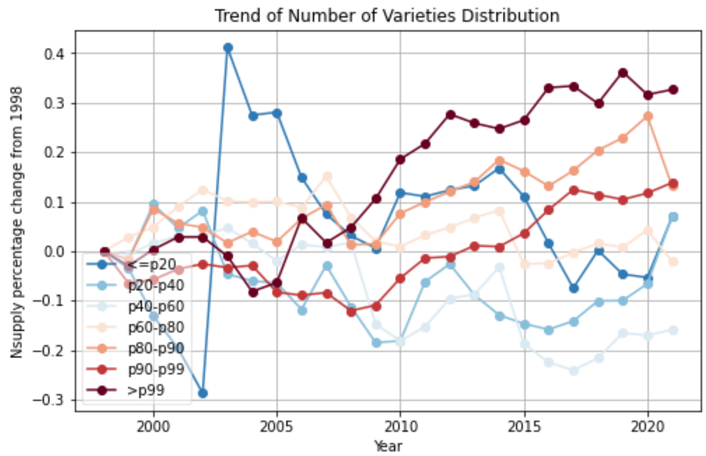

Figure 14: Number of varieties by sales percentile

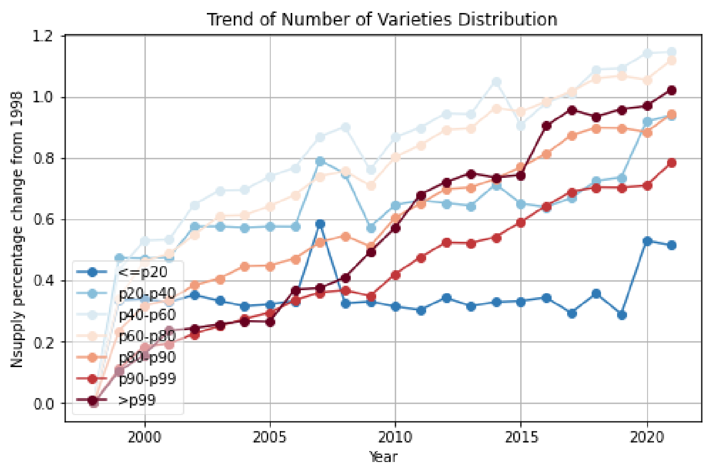

Figure 15: Number of varieties in a balanced panel

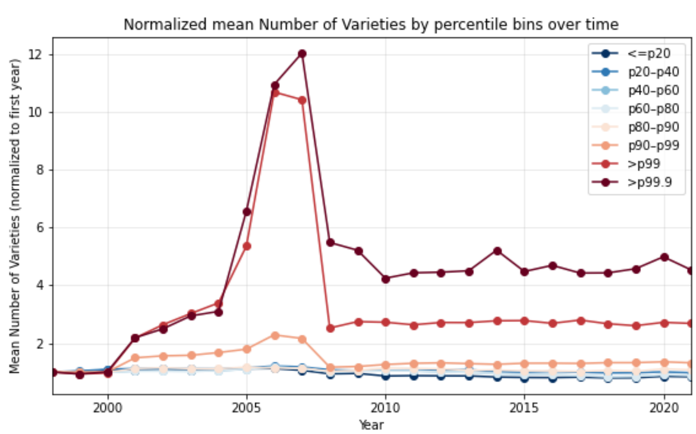

Figure 16: Number of varieties with industry by year fixed effects

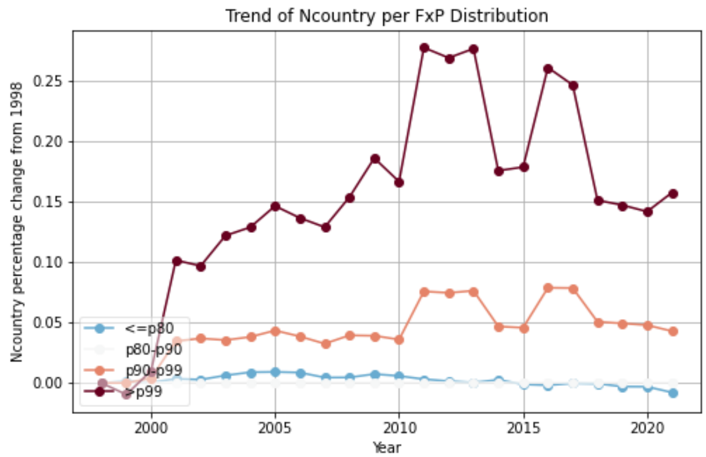

Figure 17: Number of originating countries for a firm product pair

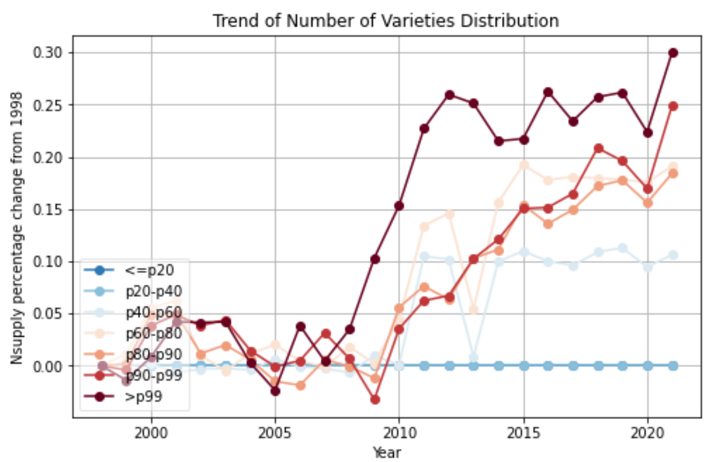

Figure 18: Number of varieties for non EU imports

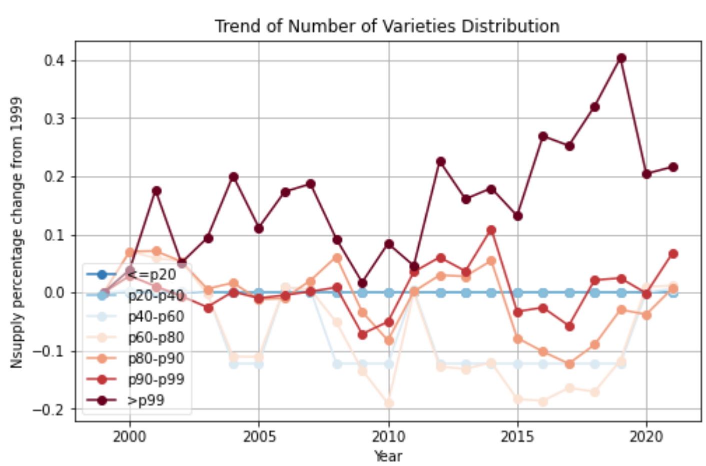

Figure 19: Number of varieties for firms without a foreign subsidiary

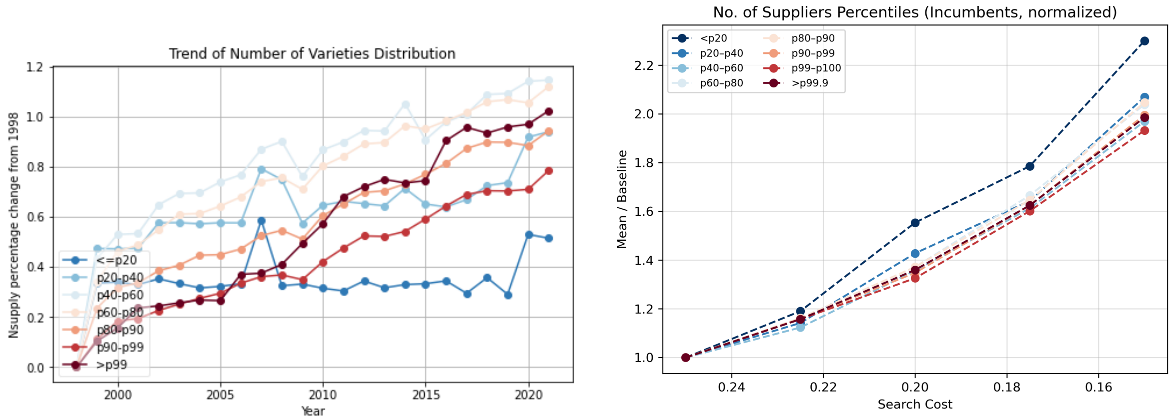

(a) Data (b) Model

Figure 20: Number of varieties dispersion in a balanced panel

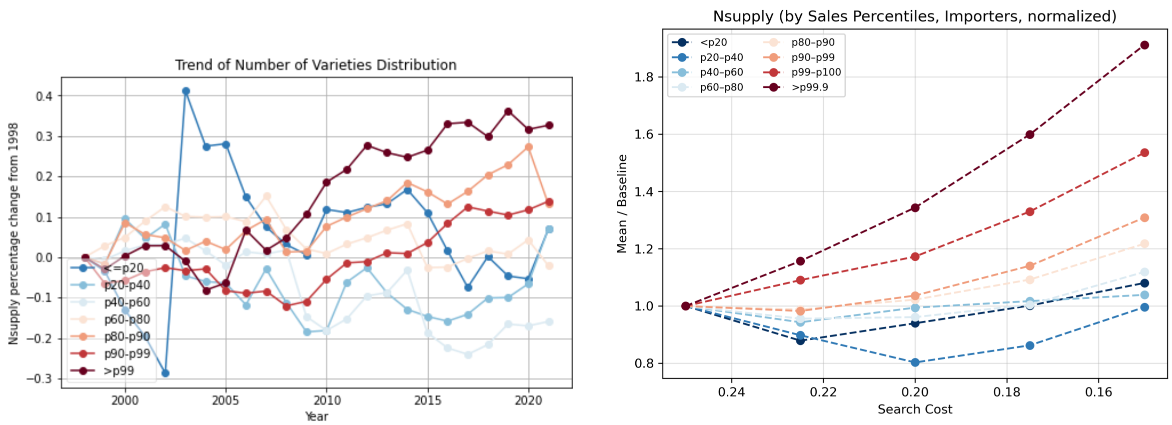

(a) Data (b) Model

Figure 21: Number of varieties dispersion by sales percentile

A.3 Robustness Tests on Sales Percentiles Evolution

There are alternative explanations for the widening firm size distribution. One possibility is that the composition of importing firms has changed for reasons unrelated to input market efficiency, such as structural shifts or sector-specific technological shocks. To address this, I conduct four robustness exercises that track the evolution of the firm size distribution across different subsamples.

First, I examine a balanced panel of firms that appear in all periods of the dataset. In this subsample, the largest firms (at the 99th percentile) exhibit the fastest growth, while other firms show similar and comparatively stable growth rates. This pattern mirrors the dispersion observed in the number of imported varieties, where the upper tail of the distribution pulls away from the rest.

Second, I demean firm sales by the industry-year mean to account for potential changes in industry composition. Although demeaning introduces some year-to-year volatility, the underlying dispersion pattern remains evident.

Third, I restrict the sample to firms without foreign subsidiaries and that are not owned by foreign parents. This ensures that the results are not driven by intra-firm trade. The dispersion pattern persists in this arm's-length trade subsample.

Finally, I focus on imports from non-EU countries. Since within-EU imports are subject to higher reporting thresholds, there is concern that measurement errors could bias the results toward larger firms. Repeating the exercise on extra-EU imports yields the same qualitative findings.

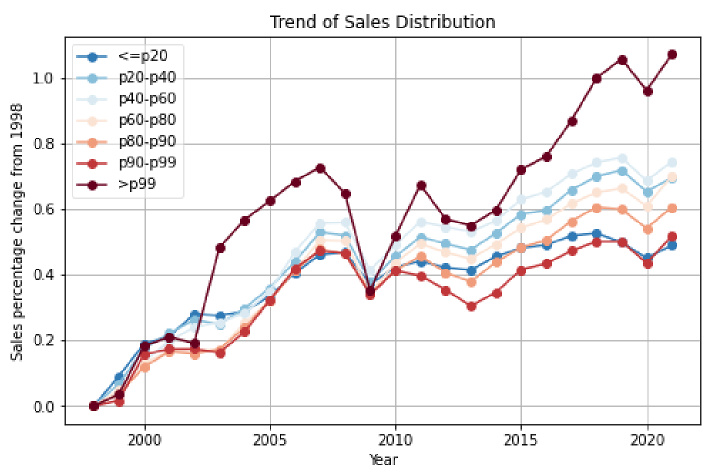

Figure 22: Sales dispersion in a balanced panel

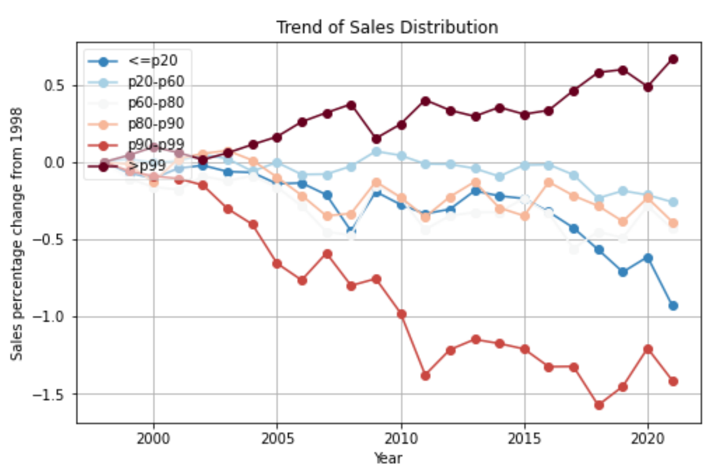

Figure 23: Sales dispersion for all manufacturing firms

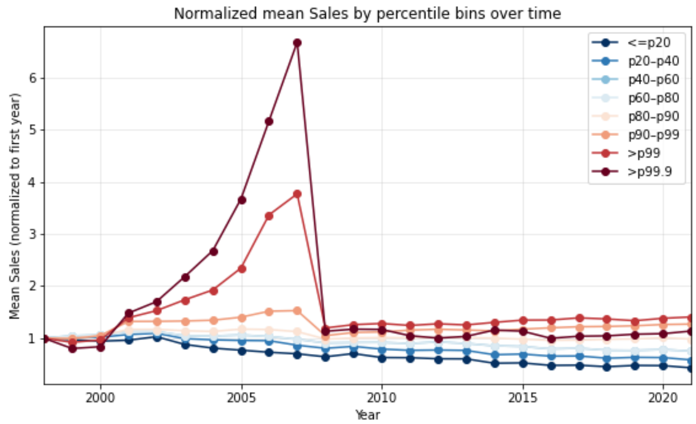

Figure 24: Sales dispersion with industry by year fixed effects

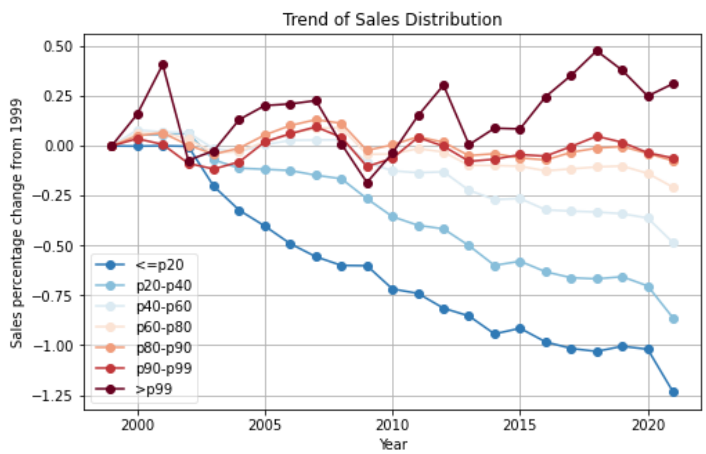

Figure 25: Sales dispersion for firms without a foreign subsidiary

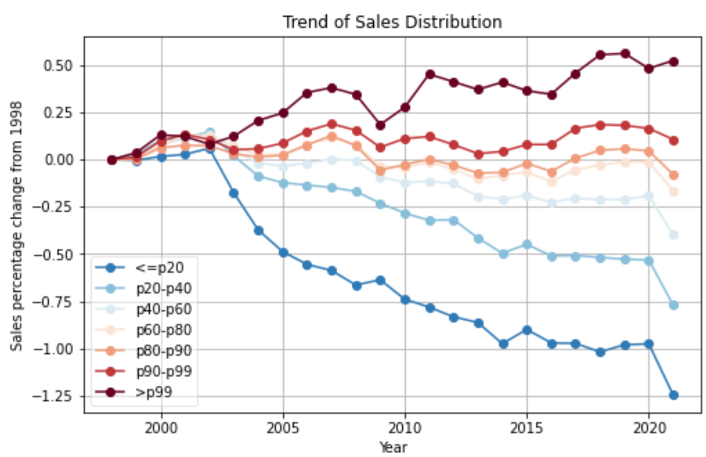

Figure 26: Sales dispersion for firms that import only from outside the EU

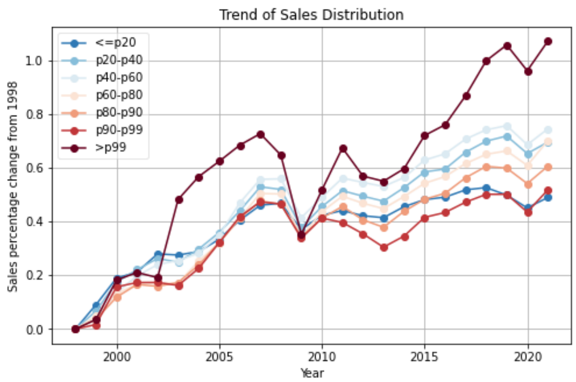

(a) Data

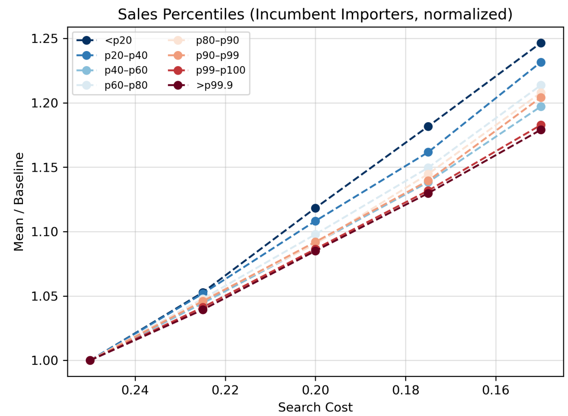

(b) Model

Figure 27: Sales dispersion in a balanced panel

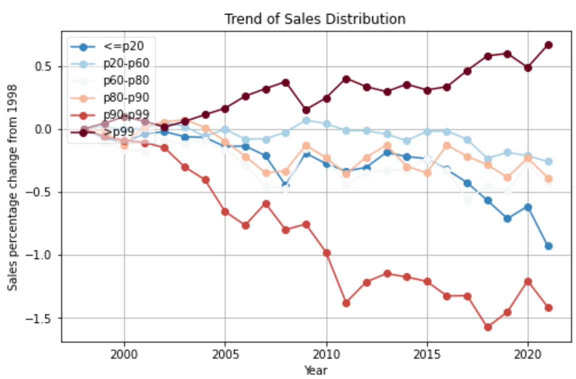

(a) Data

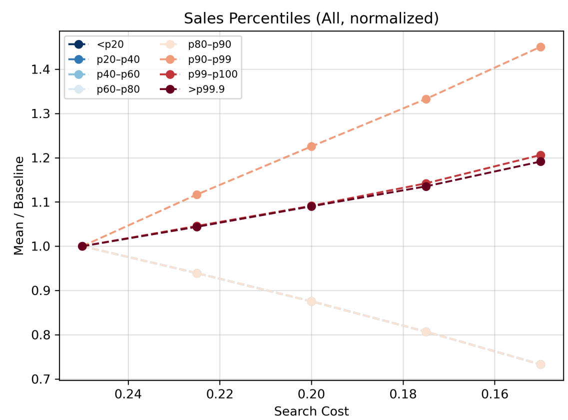

(b) Model

Figure 28: Sales dispersion for all firms

A.4 Robustness Tests on Price and Size

A.4.1 Regressing size and quantity together

Table 12: Regression results on input price

	Coefficient	(Std. error)
log number of workers	0.0048^{***}	(0.0017)
Input quantity	-0.2790^{***}	(0.0015)
Observations	172,602	
R^2	0.2838	
Industry fixed effects	Yes	
Time fixed effects	Yes	

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

When employment and input quantity are included in the same regression, the coefficient on log employment becomes small and slightly positive (0.0048), while the coefficient on input quantity remains strongly negative (−0.279). This pattern does not overturn the main result. Because employment and input quantity are highly correlated, the regression attributes most of the price–size relationship to input quantity. The small positive employment coefficient reflects residual differences in firm composition—such as labor intensity—after accounting for sourcing scale.

A.4.2 Alternative measures and other samples

Table 13: Regression results on input price: alternative measures

	Full	non EU	Arm length
TFPQ	-0.056^{***}	-0.042^{***}	-0.071^{***}
	(0.003)	(0.004)	(0.005)
TFPR	-0.336^{***}	-0.356^{***}	-0.363^{***}
	(0.007)	(0.008)	(0.008)
Output per worker	-0.234^{***}	-0.247^{***}	-0.295^{***}
	(0.013)	(0.016)	(0.019)
Output per wage	-0.150^{***}	-0.162^{***}	-0.163^{***}
	(0.006)	(0.006)	(0.007)
Value added per worker	-0.191^{***}	-0.216^{***}	-0.226^{***}
	(0.012)	(0.015)	(0.022)
Sales	-0.123^{***}	-0.125^{***}	-0.129^{***}
	(0.001)	(0.001)	(0.002)
Employment	-0.166^{***}	-0.173^{***}	-0.113^{***}
	(0.002)	(0.002)	(0.004)
Total assets	-0.105^{***}	-0.107^{***}	-0.119^{***}
	(0.001)	(0.001)	(0.001)

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

A.4.3 Realized price and quantity relationship on transaction level

The negative correlation between price and quantity is also observed at the firm–variety–year level. This pattern can arise for several reasons, including non-linear pricing, market power, or heterogeneity in product quality within a variety—firms tend to import larger quantities of lower-priced goods classified under the same product code. It is nonetheless complementary evidence for the claim in the main text: more productive firms buy more in expectation, and therefore find more efficient suppliers and pay lower unit prices.

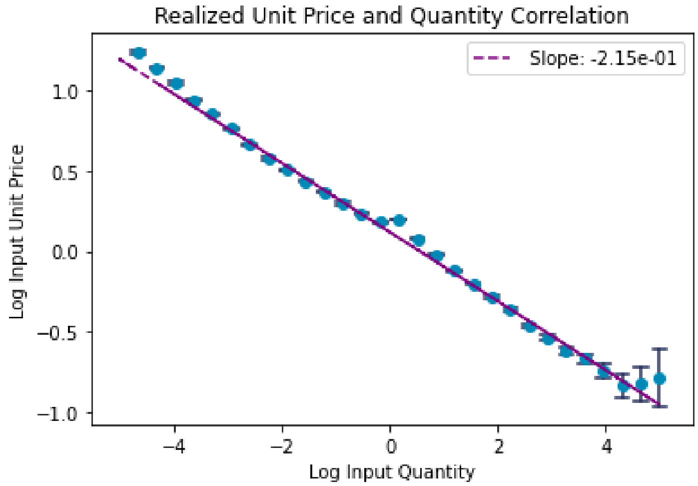

A major concern is that measurement error in quantity can mechanically create a correlation between price and quantity (Deaton, 1988), since I use value and quantity to back out the unit price. In the dataset I use, around 25% of observations have both quantity measures. I assume it is unlikely that both are subject to measurement error at the same time. Under this assumption, I compute the unit price with one measure and regress it on the other as a robustness check against such error. The shape and the coefficient change slightly in this exercise, but the negative correlation that matters is robust to measurement error.

A.5 Robustness Tests for Fiber Rollout

According to of Enterprise and Innovation (2016), the Swedish government adopted a nationwide internet strategy in 2009. Its goal is to promote fast internet for households, businesses and public services. Although profitability concerns mean that rural areas are connected more slowly, the stated aim is to close the gap between rural areas and cities over time. The treatment should therefore not be expected to relate directly to trade.

A.5.1 Municipality fiber map

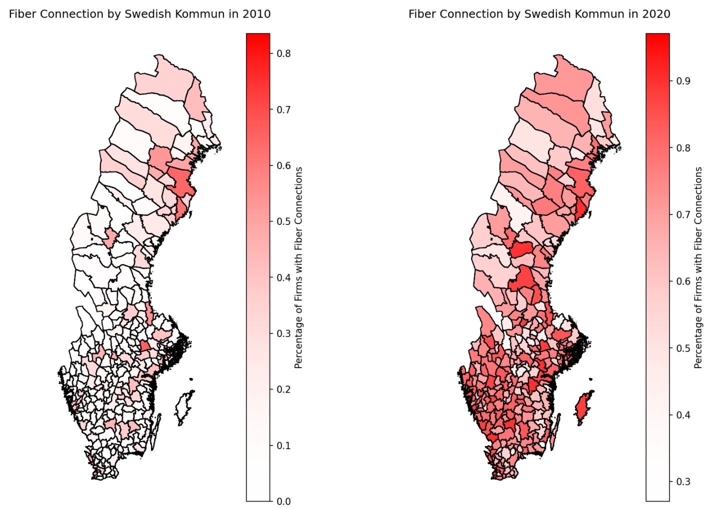

Figure 29: Municipality level fiber internet rollout

Figure 29 shows the share of firms with fiber broadband connections across Swedish kommuner in 2010 and 2020. In 2010, fiber connectivity was limited across the whole country, with relatively higher adoption in a few urban and northern regions. By 2020, coverage had expanded substantially across the country, with most municipalities showing fiber connection rates above 60–70 percent. The maps illustrate both the rapid diffusion of broadband infrastructure and the remaining regional disparities in adoption.

A.5.2 Robustness to Pretrend

This appendix reports a binary pretrend robustness that complements the continuous coverage specification in the main text. I set $Event_{m,t} = 1$ when municipal fiber coverage exceeds 50 percent, and define *PreEvent* as one in the three years preceding the event. The dependent variable is measured in levels of supplier network size rather than in log or percentage changes.

The big firm indicator equals one for firms with sales above the contemporaneous industry mean.

$$
\begin{aligned}
\text{Varieties}_{f,t} = {} & \beta_1 \, \text{Event}_{m,t} + \beta_2 \Big[\text{Event}_{m,t} \times \text{BigFirm}_{f,t} \Big] \\
& + \gamma_1 \, \text{PreEvent}_{m,t} + \gamma_2 \Big[\text{PreEvent}_{m,t} \times \text{BigFirm}_{f,t} \Big] \\
& + \text{FE}_t + \text{FE}_i + \text{FE}_m + \varepsilon_{f,t}.
\end{aligned}
$$

Table 14: Connectivity and supplier network size with pre-event control (sales based big)

	(1)	(2)	(3)	(4)
Treated	1.9191***	2.3428***	0.0404	0.1406
	(0.5809)	(0.6420)	(0.0856)	(0.1879)
Treated × BigFirm			9.328***	11.527***
			(2.576)	(2.871)
PreEvent		0.6608		0.1541
		(0.4473)		(0.1823)
PreEvent × BigFirm				3.4857**
				(1.642)
BigFirm (main effect)			21.934***	19.736***
			(1.053)	(1.228)
Observations	349,775	349,775	349,775	349,775
R^2	0.0237	0.0237	0.0675	0.0675
Industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Municipality FE	Yes	Yes	Yes	Yes

Notes: Dependent variable is the level of imported input varieties for firm f in year t. "Treated" equals one when municipal fiber coverage exceeds 50 percent. "PreEvent" equals one in the three pre-event years. "BigFirm" is an indicator equal to one for firms with sales above the contemporaneous industry mean. Standard errors clustered by municipality in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

The estimates indicate that supplier network size expands more strongly for large firms following fiber rollout, while pretrend coefficients are small and statistically insignificant.

A.5.3 Robustness Using Growth Rate Specification

As a complementary exercise, I estimate the relationship between connectivity and the *percentage change* in a firm's imported varieties. This outcome allows for comparability across firms of different sizes, although the theoretical predictions refer to levels rather than growth rates. The specification parallels the baseline regression but uses the change in imported varieties relative to the previous year as the dependent variable:

$$\begin{aligned}\%\Delta\text{Varieties}_{f,t} = \beta_1\,\text{FiberCoverage}_{m,t} + \beta_2\Big[\text{FiberCoverage}_{m,t} \times \text{BigFirm}_{f,t}\Big] \\ + \text{FE}_t + \text{FE}_i + \text{FE}_m + \varepsilon_{f,t}.\end{aligned}$$

The big firm indicator equals one for firms with sales above the contemporaneous industry mean.

Table 15: Connectivity and the growth rate of imported varieties

	(1)	(2)
Fiber coverage	0.1485***	0.1064***
	(0.0160)	(0.0184)
Fiber × BigFirm	—	0.1817***
		(0.0325)
Observations	112,022	112,022
R^2	0.0013	0.0020
Industry FE	Yes	Yes
Year FE	Yes	Yes
Municipality FE	Yes	Yes

Notes: Dependent variable is the percentage change in a firm's number of imported input varieties. "BigFirm" is an indicator equal to one for firms with sales above the contemporaneous industry mean. Standard errors clustered by municipality in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

The estimates show that variety growth is faster in municipalities with higher fiber coverage, and the effect is substantially stronger for large firms. While the R^2 values are small, consistent with firm-level heterogeneity and municipality-level aggregation, the pattern supports the idea that reduced search costs expand supplier networks more strongly among large firms.

A.6 Other Empirical Specifications

A.6.1 Quality-adjusted TFPQ

In the main part of the paper, I used a direct measure of size and productivity. That approach has two potential drawbacks. First, it selects a small subsample in which output prices are observed. These are usually larger firms, so the results may be biased. Second, a direct output price measure does not adjust for quality differences across firms.

I define the revenue productivity the same way as above, but then I estimate the quality-adjusted TFPQ by:

$$\text{TFPQ}_f = \frac{Q_f}{p_f}\text{TFPR}_f$$

, where p is the price and Q is the quality of the firm's production. Unfortunately, quality Q is generally not observable. To obtain an estimate of Q, I have to rely on the theory. Assuming the monopolistic competition in final good demand, the household optimization gives:

$$\frac{p_i y_i}{p_j y_j} = (\frac{p_i/Q_i}{p_j/Q_j})^{1-\epsilon}$$

where sales py are observed in the data. I can then back out $\frac{p}{Q}$ by applying the standard elasticity $\epsilon = 5$. This yields an estimate of quality-adjusted TFPQ that accounts for quality without substantially reducing the sample size. I repeat the same analysis as above:

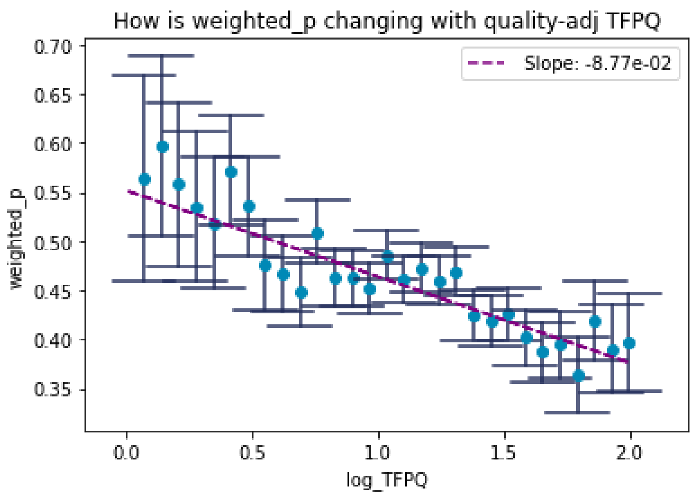

Figure 30: Relationship between quality-adjusted TFPQ and weighted prices

As predicted, the coefficient is also negative with a smaller magnitude than the non-quality adjusted version. As explained in Kugler and Verhoogen (2012), firms with higher productivity usually use higher quality inputs, which dampens the quality-adjusted coefficient.

Using a different aggregation level (industry digit), or including capital or not, also does not change the result much.

A.6.2 TFPQ and Variety

Another productivity-related empirical fact is that more productive firms buy more varieties. This is further evidence that productive firms are likely to benefit from falling search costs, because they search more. A productive firm that expects to produce more has a stronger incentive to look for efficient suppliers, and to hold several of them.

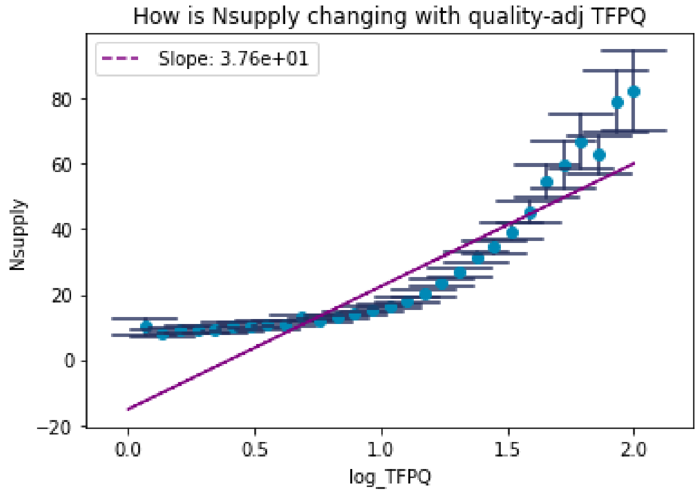

Figure 31: Relationship between TFPQ and number of imported varieties

A.6.3 Input Price and Profit

Input prices also affect a firm's profitability directly. I find a negative correlation between firm-level input prices and the profit share, even controlling for industry fixed effects. This suggests that the standard CES model with a constant markup misses important features of the intermediate goods market.

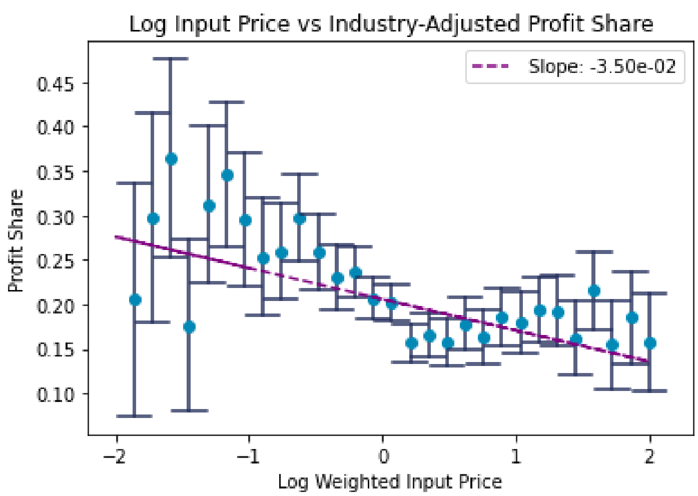

Figure 32: Relationship between input price and profit share

A.6.4 Quality Differences

Since Kugler and Verhoogen (2012) show that quality is an important channel of input price differences, I estimate how much of the dispersion in input prices quality differences can generate.

I use a reduced form approach based on the Rauch (1999) classification, which separates goods into three classes: traded on an organized exchange, traded at a referenced price, and differentiated. The first class also includes products that have an organized exchange but can be traded decentrally as well, such as bananas, wool and other basic agricultural products. These products have very little quality variation and are therefore suitable reference goods for estimating quality differentiation.

The Rauch classification is based on SITC Rev. 2 while the main dataset uses CN codes, so I link the two using the HS-SITC conversion tables provided by UNSD[12]. CN codes share their first six digits with the HS system, and I assume that all products sharing those six digits receive the same Rauch classification. I then compare the standard deviation of prices across the three groups:

Table 16: Price dispersion by Rauch classification

Category	Standard deviation
Organized exchange	0.90
Referenced price	1.18
Differentiated goods	1.44

Notes: Standard deviation of log unit prices by Rauch (1999) classification.

Although differentiated goods exhibit higher standard deviations, part of the price dispersion also persists within the "organized exchange" category. This exercise shows that quality differences alone cannot account for all of the observed variation in input prices.

The key point is that there is no systematic trend in quality differences over time. Since non-differentiated goods have minimal quality variation, I define the quality gap as the difference in price variance between differentiated and non-differentiated goods. Over the observed years, this quality gap remains largely unchanged. Using weighted averages, there is virtually no variation at all. Moreover, the share of differentiated goods in total imports is stable, ranging between 74% and 76%. Taken together, these patterns indicate that changes in product quality cannot explain the observed increase in price variance.

[12] https://unstats.un.org/unsd/classifications/Econ

B Theory Appendix

B.1 Closed Form Solution for Important Variables

In the baseline model with imported inputs, the firm's state is represented by the pair (z, X). The variables required for the value function iteration, which are the optimal price $p(z, X)$, output $y(z, X)$, and profits $\pi(z, X)$, can all be expressed as functions of this state. The problem is therefore fully characterized by the pair (z, X).

Output price and output.

$$\max_{p,y} py - w\left(\tfrac{y}{zX^{\alpha}}\right)^{\frac{1}{1-\alpha}} \quad \text{s.t.} \quad y = C\left(\tfrac{p}{P}\right)^{-\epsilon}$$

Using the constraint to eliminate y:

$$\max_{p} p^{1-\epsilon}\,\frac{C}{P^{-\epsilon}} \; - \; w\,p^{\frac{\epsilon}{\alpha-1}}\left(\frac{\frac{C}{P^{-\epsilon}}}{zX^{\alpha}}\right)^{\frac{1}{1-\alpha}}$$

FOC in p:

$$(1-\epsilon)\,\frac{C}{P^{-\epsilon}}\,p^{-\epsilon} = \frac{w\epsilon}{\alpha-1}\,p^{\frac{\epsilon-\alpha+1}{\alpha-1}}\left(\frac{\frac{C}{P^{-\epsilon}}}{zX^{\alpha}}\right)^{\frac{1}{1-\alpha}}$$

Solve for p:

$$\begin{aligned}
p^{\frac{1+\alpha(\epsilon-1)}{1-\alpha}} &= \frac{w\epsilon}{(\epsilon-1)(1-\alpha)}\left(\tfrac{C}{P^{-\epsilon}}\right)^{\frac{\alpha}{1-\alpha}}\left(zX^{\alpha}\right)^{-\frac{1}{1-\alpha}}, \\
p &= \left(\tfrac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{1-\alpha}{1+\alpha(\epsilon-1)}}\left(\tfrac{C}{P^{-\epsilon}}\right)^{\frac{\alpha}{1+\alpha(\epsilon-1)}}\left(zX^{\alpha}\right)^{-\frac{1}{1+\alpha(\epsilon-1)}}
\end{aligned}$$

Then y:

$$\begin{aligned}
y &= \frac{C}{P^{-\epsilon}}\,p^{-\epsilon} \\
&= \left(\tfrac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{-\epsilon(1-\alpha)}{1+\alpha(\epsilon-1)}}\left(\tfrac{C}{P^{-\epsilon}}\right)^{\frac{1-\alpha}{1+\alpha(\epsilon-1)}}\left(zX^{\alpha}\right)^{\frac{\epsilon}{1+\alpha(\epsilon-1)}}
\end{aligned}$$

Profit function or not search value V^{NS}

$$
\begin{aligned}
\pi(z,X) &= py \; - \; w\bigg(\frac{y}{zX^{\alpha}}\bigg)^{\frac{1}{1-\alpha}} \\
&= \frac{C}{P^{-\epsilon}}\, p^{1-\epsilon} \; - \; w\left(\frac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{-\epsilon}{1+\alpha(\epsilon-1)}}\left(\frac{C}{P^{-\epsilon}}\right)^{\frac{1}{1+\alpha(\epsilon-1)}} \\
&\quad \times \left(zX^{\alpha}\right)^{\frac{\epsilon-1-\alpha(\epsilon-1)}{1+\alpha(\epsilon-1)}\cdot\frac{1}{1-\alpha}} \\
&= \left\{\left(\frac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{(1-\alpha)(1-\epsilon)}{1+\alpha(\epsilon-1)}} - w\left(\frac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{-\epsilon}{1+\alpha(\epsilon-1)}}\right\} \\
&\quad \times \left(\frac{C}{P^{-\epsilon}}\right)^{\frac{1}{1+\alpha(\epsilon-1)}}\left(zX^{\alpha}\right)^{\frac{\epsilon-1}{1+\alpha(\epsilon-1)}} \\
&= \left[\frac{1+\alpha(\epsilon-1)}{\epsilon}\right]\left(\frac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{(1-\alpha)(1-\epsilon)}{1+\alpha(\epsilon-1)}}\left(\frac{C}{P^{-\epsilon}}\right)^{\frac{1}{1+\alpha(\epsilon-1)}} \\
&\quad \times \left(zX^{\alpha}\right)^{\frac{\epsilon-1}{1+\alpha(\epsilon-1)}}
\end{aligned}
$$

Note that the profit function π is also the not-search value V^{NS}.

Labor cost.

$$
\begin{aligned}
w\bigg(\frac{y}{zX^{\alpha}}\bigg)^{\frac{1}{1-\alpha}} &= \left(\frac{C}{P^{-\epsilon}}\right)^{\frac{1}{1+\alpha(\epsilon-1)}} w\left(\frac{w\epsilon}{(\epsilon-1)(1-\alpha)}\right)^{\frac{-\epsilon}{1+\alpha(\epsilon-1)}}\left(zX^{\alpha}\right)^{\frac{\epsilon-1}{1+\alpha(\epsilon-1)}} \\
&= \frac{(\epsilon-1)(1-\alpha)}{\epsilon}\, py \quad \text{or} \quad \left[\frac{(\epsilon-1)(1-\alpha)}{1+\alpha(\epsilon-1)}\right]\pi
\end{aligned}
$$

B.2 Simplified Benchmark - Single Supplier Case

This is another simplified model, where firms can have only one supplier (variety), there is no quality differentiation, and both upstream and downstream productivity follow a $U(0,1)$ distribution. It resembles the random job search model of McCall (1970). With this class of model I can solve the optimal stopping problem for the reservation productivities $z_u^R(z)$ and z^R. The fixed point solution then shows how the search cost κ affects the firm size distribution.

The reservation supplier productivity is solved by 2 equations. The first equation is the definition of the outside option D:

$$
D = \int_{z_u^R}^{\bar{z_u}} [\pi(z_u) - T(z_u)]dF(z_u) + \int_{\underline{z}_u}^{z_u^R} D dF(z_u) - w\kappa
$$

which consists of 3 terms: the first is expected profit given the probability of meeting a satisfactory supplier ($z_u \geq z_u^R$), the second is falling back on the outside option if the supplier is worse than the reservation productivity, and the last is minus the search cost.

The second equation is the indifference condition: any buyer must be indifferent between the profit from meeting the reservation supplier and taking the outside option:

$$\pi(z, z_u^R) - T(z, z_u^R) = D$$

Combining the equations gives an expression for solving z_u^R for each z:

$$\frac{\kappa}{Kz^{\epsilon-1}} = [1 + \alpha(\epsilon - 1)z_u^{R[\alpha(\epsilon-1)+1]} - (\alpha(\epsilon-1)+1)z_u^{R\alpha(\epsilon-1)}]$$

This expression provides a relationship between z_u^R and κ for any firm with productivity z. Solving this fixed point formula gives:

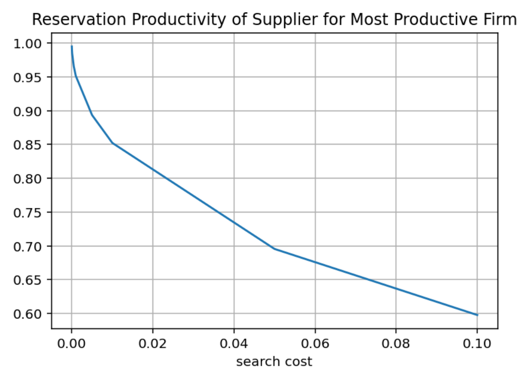

This shows that when the search cost decreases, the reservation productivity, given z, increases. In other words, when the search cost decreases, the expected productivity of the supplier increases. Matching with a more productive supplier means producing more. Therefore, downstream firms produce more when the search cost decreases. In this restrictive model the search cost also functions as an expansion cost.

On top of that, I can show that the lowest productivity z^R where this downstream firm still remains active can be calculated by solving:

$$\frac{\kappa}{1 - F[z_u^R(z^R)]} = \int_{z_u^R(z^R)}^{\bar{z_u}} [\pi(z_u, z^R) - T(z_u, z^R)] dF(z_u)$$

Holding everything but the search cost κ constant, z^R increases with κ:

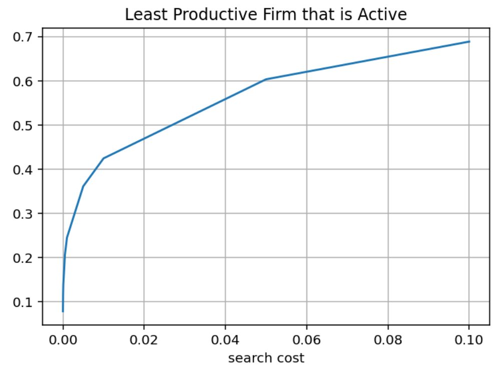

It shows that when the search cost decreases, less productive firms can also enter the intermediate good market and produce. The search cost therefore acts as an entry barrier.

B.3 Sequential Search with Efficient Bargaining

This appendix solves the sequential version of the buyer's problem with identical upstream productivity and constant search cost as described in A Simplified Case: Price Dispersion. Upstream unit cost is fixed at $c = 1/\bar{z}_u$. At any state (z, X) the buyer meets a new supplier and bargains efficiently over x_i. After contracting, total input becomes

$$X_{\text{new}} = (X^\rho + x_i^\rho)^{1/\rho}.$$

Define

$$A \equiv \frac{\epsilon - 1}{1 + \alpha(\epsilon - 1)}, \qquad K_1 \equiv [1 + \alpha(\epsilon - 1)]\, \epsilon^{-\frac{\epsilon}{1+\alpha(\epsilon-1)}} \left(\frac{w}{(\epsilon-1)(1-\alpha)}\right)^{\frac{(1-\alpha)(1-\epsilon)}{1+\alpha(\epsilon-1)}} \left(\frac{C}{P^{-\epsilon}}\right)^{\frac{1}{1+\alpha(\epsilon-1)}}.$$

With this notation $K_1[zX^\alpha]^A$ is the closed-form profit function $\pi(z, X)$ derived in Closed Form Solution for Important Variables.

In the main text the value function $V(z, X)$ is sequential and includes both the direct production benefit and the continuation value from future supplier matches. For the derivation it is convenient to write

$$V(z, X) = K_1[zX^\alpha]^A + W(z, X),$$

so that

$$V(z, X_{\text{new}}) - V(z, X) = K_1\Big([zX_{\text{new}}^\alpha]^A - [zX^\alpha]^A\Big) + \Big[W(z, X_{\text{new}}) - W(z, X)\Big].$$

Per match bargaining problem At (z, X) the efficient surplus from contracting x_i is

$$\max_{x_i \geq 0} \underbrace{K_1\Big[z(X^\rho + x_i^\rho)^{\alpha/\rho}\Big]^A - K_1\Big[zX^\alpha\Big]^A}_{\text{direct benefit}} + \underbrace{W(z, (X^\rho + x_i^\rho)^{1/\rho}) - W(z, X)}_{\text{continuation effect}} - c\,x_i.$$

Let $x_i^*(z, X)$ denote the maximizer of this problem and write $x^* \equiv x_i^*(z, X^*)$ at the stopping boundary introduced below. The buyer's surplus from the match is

$$g(z, X) = \theta\Big\{K_1\Big[z(X^\rho + x_i^\rho)^{\alpha/\rho}\Big]^A - K_1\Big[zX^\alpha\Big]^A + W(z, (X^\rho + x_i^\rho)^{1/\rho}) - W(z, X) - c\,x_i\Big\}.$$

First order condition Let $X_{\text{new}} = (X^\rho + x_i^\rho)^{1/\rho}$. The first order condition with respect to x_i is

$$\Big[\alpha A K_1 z^A (X^\rho + x_i^\rho)^{\frac{\alpha A}{\rho} - 1} + W_2(z, X_{\text{new}})(X^\rho + x_i^\rho)^{\frac{1}{\rho} - 1}\Big] x_i^{\rho - 1} = c,$$

or, equivalently,

$$\Big[\alpha A K_1 z^A X_{\text{new}}^{\alpha A - 1} + W_2(z, X_{\text{new}})\Big] x_i^{\rho-1} X_{\text{new}}^{1-\rho} = c.$$

Search rule and stopping boundary The gross return to searching rather than stopping is $G(z, X) = W(z, X) + g(z, X)$, the option value already in hand plus the surplus from the next match. Search continues while $G(z, X) > w\kappa$. The endogenous stopping boundary $X^*(z)$ satisfies

$$G(z, X^*) = w\kappa,$$

which is $g(z, X^*) = w\kappa$ because $W(z, X^*) = 0$. At X^* the continuation effect from the marginal match vanishes: every stock above X^* is in the stopping region and $x^* > 0$ puts X_{new} there, so $W_2(z, X_{\text{new}}) = 0$, and the first order condition reduces to

$$\alpha A K_1 z^A (X^{*\rho} + x_i^{*\rho})^{\frac{\alpha A}{\rho} - 1} x_i^{*\rho - 1} = c.$$

Solving in ratios Define $r \equiv x^*/X^*$ and note $X_{\text{new}} = X^*(1 + r^\rho)^{1/\rho}$. The first order condition implies

$$X^{*\alpha A - 1} = \frac{c}{\alpha A K_1 z^A (1 + r^\rho)^{\frac{\alpha A - \rho}{\rho}} r^{\rho - 1}}, \qquad X^* = \left[\frac{c}{\alpha A K_1 z^A (1 + r^\rho)^{\frac{\alpha A - \rho}{\rho}} r^{\rho - 1}}\right]^{\frac{1}{\alpha A - 1}}.$$

At the boundary, the surplus condition becomes

$$K_1\Big[z(X^{*\rho}+x_i^{*\rho})^{\alpha/\rho}\Big]^A - K_1(zX^{*\alpha})^A - c\,x^* = \frac{w\kappa}{\theta},$$

which, after substituting $r = x^*/X^*$ and the expression for $X^{*\alpha A-1}$, reduces to

$$X^* c\left[\frac{(1+r^\rho)-(1+r^\rho)^{1-\frac{\alpha A}{\rho}}}{\alpha A\, r^{\rho-1}} - r\right] = \frac{w\kappa}{\theta}.$$

This condition and the first order condition are two equations in the two unknowns X^* and r. Solving them yields $X^*(z)$ and $x^*(z) = r\,X^*(z)$. The model can then be solved recursively by backward induction starting from the stopping boundary.

Contract quantity and buyer productivity For a given existing input stock X, the efficient contract quantity satisfies

$$V_X(z, X_{\text{new}})x_i^{\rho-1}X_{\text{new}}^{1-\rho} = c, \qquad X_{\text{new}} = (X^\rho + x_i^\rho)^{1/\rho}.$$

Implicit differentiation gives

$$\frac{\partial x_i^*}{\partial z} = -\frac{V_{Xz}(z, X_{\text{new}})x_i^{\rho-1}X_{\text{new}}^{1-\rho}}{\frac{\partial^2}{\partial x_i^2}\left[V(z, X_{\text{new}}) - cx_i\right]}.$$

Therefore, if the match surplus is strictly concave in x_i and $V_{Xz}(z, X_{\text{new}}) > 0$, then

$$\frac{\partial x_i^*(z, X)}{\partial z} > 0.$$

B.4 Concavity of the Bargaining Problem

Given all possible x, I can use the first-order condition (FOC) to find T, because the following function is strictly concave (as long as it is defined). To simplify notation, I define

$$V(x) \coloneqq V(z, X_N) - V(z, X), \qquad K(x) \coloneqq \frac{w}{z_u}x.$$

Lemma 1. *Fix x with $V(x) > K(x)$. Then $\log B(x,\cdot)$ is strictly concave on $(K(x), V(x))$, and $B(x,\cdot)$ has a unique maximizer $T(x)$.*

Proof. Consider

$$B(\delta T_1 + (1-\delta)T_2) = \Big(V(x) - \delta T_1 - (1-\delta)T_2\Big)^{\theta} \Big(\delta T_1 + (1-\delta)T_2 - K(x)\Big)^{1-\theta}.$$

Taking log of both sides,

$$\begin{aligned}\log B(\delta T_1 + (1-\delta)T_2) = {} & \theta \log\Big(V(x) - \delta T_1 - (1-\delta)T_2\Big) \\ & + (1-\theta)\log\Big(\delta T_1 + (1-\delta)T_2 - K(x)\Big).\end{aligned}$$

Because the logarithm is strictly concave, for any $\delta \in (0,1)$,

$$\begin{aligned}\log\Big(V(x) - \delta T_1 - (1-\delta)T_2\Big) &\geq \delta \log\Big(V(x) - T_1\Big) + (1-\delta)\log\Big(V(x) - T_2\Big), \\ \log\Big(\delta T_1 + (1-\delta)T_2 - K(x)\Big) &\geq \delta \log\Big(T_1 - K(x)\Big) + (1-\delta)\log\Big(T_2 - K(x)\Big).\end{aligned}$$

Substituting these inequalities into the expression for $\log B$ and grouping terms gives

$$\begin{aligned}\log B(\delta T_1 + (1-\delta)T_2) &\geq \delta\Big[\theta \log(V(x) - T_1) + (1-\theta)\log(T_1 - K(x))\Big] \\ &\quad + (1-\delta)\Big[\theta \log(V(x) - T_2) + (1-\theta)\log(T_2 - K(x))\Big] \\ &= \delta \log B(T_1) + (1-\delta)\log B(T_2).\end{aligned}$$

Hence $\log B(T)$ is strictly concave on $(K(x), V(x))$. Since log is strictly increasing, $B(T)$ is strictly log-concave. As long as $V(x) \geq K(x)$, the maximizer of $B(T)$ exists and is unique. □

Lemma 2. *Any maximizer (x^*, T^*) of B satisfies $T^* = T(x^*)$.*

Proof. Let $T(x^o) = \arg\max_T B(x^o, T)$, which is unique by Lemma 1. Assume (x^*, T^*) is the maximizer of B such that $B(x^*, T^*) \geq B(x', T')$ for all (x', T'). If $T^* \neq T(x^*)$, then $B(x^*, T^*) \geq B(x^*, T(x^*))$, which contradicts the uniqueness of $T(x^*)$. Therefore, the solution is unique. □

We can therefore define a function *Bargain* that gives the best value given x, and then find $\arg\max_x$ Bargain.

I can then obtain *Bargain* (the one-dimensional version of the bargaining problem) by substituting the FOC into B:

$$\begin{aligned}
\text{Bargain}(x) &= (-T + V(z, X_N) - V(z, X))^{\theta}(T - w\tfrac{x}{z_u})^{1-\theta} \\
&= \Big(-(1-\theta)V(x) - \theta K(x) + V(x)\Big)^{\theta}\Big((1-\theta)V(x) + \theta K(x) - K(x)\Big)^{1-\theta} \\
&= \Big(\theta V(x) - \theta K(x)\Big)^{\theta}\Big((1-\theta)V(x) - (1-\theta)K(x)\Big)^{1-\theta} \\
&= \theta^{\theta}(1-\theta)^{1-\theta}\Big(V(x) - K(x)\Big) \\
&= \theta^{\theta}(1-\theta)^{1-\theta}\Big(V(z, X^N Q^N, n+1) - V(z, XQ, n+1) - w\tfrac{x}{z_u}\Big).
\end{aligned}$$

B.5 Computational Details

This appendix summarizes the numerical implementation used in Section 5. Calibrated parameters are reported in Tables 2 and 3.

Grids. Buyer productivity z and supplier productivity z_u use 20- and 40-node grids, respectively. Both are lognormal with $\mu = 0$ and $\sigma = 0.235$, truncated at the 0.01 and 99.99 percentiles, with support [0.417, 2.397]. The endogenous state X and choice x use 6,000 and 375 log-spaced nodes, respectively, with zero included as the first node. The fine X grid captures the discontinuous search policy and the skewed distribution of X.

Upper boundary. The ceilings for x and X scale with $z_{\max}$ and $z_{u,\max}$: $X_{\max} = 8{,}618$ and $x_{\max} = 86.2$. I also set a maximum number of searches per firm, 6,000, to prevent infinite loops. The boundary is set deliberately high so that, given the current productivity setting, the upper bound is never reached. I reject simulations in which a firm reaches either grid's top two nodes or exhausts the search cap.

Implementation. *Precomputed interpolation maps.* For each pair of a state node X and a candidate input x, the implied $X_{\text{new}} = (X^{\rho} + x^{\rho})^{1/\rho}$ is located on the log-spaced X grid once, before iteration begins, and stored as a lower-neighbor index and an interpolation weight. The map depends only on ρ and the two grids, so the $6{,}000 \times 375$ lookups are paid once rather than at every continuation value thereafter. *Parallelization.* The operator is block diagonal in z, so the 20 rows are updated concurrently with thread-local buffers. Work is assigned dynamically rather than in fixed blocks, because rows converge and retire at different times.

Convergence. Convergence is assessed row by row: a row retires once its own sup-norm residual falls below tolerance on two consecutive iterations, and the inner loop ends when no rows remain active. The relaxation weight is bounded to [0.95, 2]. Value-function iteration uses a loose tolerance when the aggregate demand term is far from market clearing and 10^{-3} near convergence. The outer loop solves for CP^{ϵ} using Brent's method with relative tolerance 10^{-4}. Each $\kappa \in \{0.25, 0.225, 0.20, 0.175, 0.15\}$ is solved as a separate stationary equilibrium.