

NBER WORKING PAPER SERIES

INPUT PRICE DISPERSION ACROSS BUYERS

Ariel Burstein
Javier Cravino
Marco Rojas

Working Paper 33128
<http://www.nber.org/papers/w33128>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2024, Revised June 2026

We thank Ken Kikkawa, John Haltiwanger, Federico Huneeus, Ernesto Pastén, Chad Syverson, Nicholas Trachter, Jonathan Vogel, and seminar and conference participants at various institutions for valuable comments and suggestions. The views expressed are those of the authors and do not necessarily represent the views of the Central Bank of Chile, its board members, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Ariel Burstein, Javier Cravino, and Marco Rojas. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Input Price Dispersion Across Buyers
Ariel Burstein, Javier Cravino, and Marco Rojas
NBER Working Paper No. 33128
November 2024, Revised June 2026
JEL No. D22, D24, D4, E02, O4

ABSTRACT

We use a comprehensive dataset of electronic invoices from Chilean firms to document new facts about price dispersion across buyers of the same manufactured intermediate input and assess its implications for aggregate productivity. Price dispersion is pervasive: more than half of manufacturing sales involve products sold to multiple buyers within the same month, with an average interdecile range of prices across buyers of 0.16 log points. The dominant source of this dispersion is match-specific: buyer-seller pair identity absorbs more than 40 percent of the variation in price gaps, with observable buyer-seller characteristics — including size, relationship length, and transaction volume — jointly explaining up to 10 percent. Price gaps vary little with geographic distance or payment method (cash vs. credit), suggesting that they primarily reflect markup differences rather than marginal cost differences. Guided by these facts, we use a production network model calibrated to the Chilean data to quantify the aggregate efficiency gains from eliminating markup dispersion in manufactured intermediate inputs. We find that markup heterogeneity across buyers of the same product accounts for 30 to 64 percent of the total productivity gains from fully equalizing markups across both products and buyers.

Ariel Burstein
University of California, Los Angeles
Department of Economics
and NBER
arielb@econ.ucla.edu

Marco Rojas
Central Bank of Chile
mrojaso@bcentral.cl

Javier Cravino
University of Michigan
Department of Economics
and NBER
jcravino@umich.edu

1 Introduction

To what extent do firms pay different prices for the same intermediate input? Which observable buyer-seller characteristics account for the variation in prices across buyers? Does input price dispersion across buyers generate misallocation in domestic production networks, thereby reducing aggregate productivity? Answering these questions is important for guiding models of firm-to-firm trade and for understanding sources of inefficiencies in domestic supply chains. Yet, existing evidence on price dispersion for intermediate inputs is scarce and largely limited to products within narrow industries or to international trade transactions.¹

This paper provides the first comprehensive evidence on price dispersion across buyers of intermediate inputs and quantifies its aggregate productivity implications. We use transaction-level data from electronic invoices issued by Chilean manufacturing firms between June 2021 and December 2023. These invoices include seller and buyer identities, transaction values and quantities, and detailed product codes and descriptions from firms' billing systems. This information allows us to identify products at a much finer level than standard industry or customs classifications, and to track how the price of the same product sold by the same seller varies across buyers within the same date. The set of uniquely identified products accounts for nearly half of all Chilean firm-to-firm manufacturing sales, and the industry composition and price changes in our data closely resemble those in the Chilean producer price index.

Our first empirical contribution is to document that price dispersion across buyers of the same intermediate input is pervasive, persistent, and unlikely to be driven by cost differences. More than half of manufacturing sales involve products sold to multiple buyers within the same month, with an average interdecile range of prices across buyers of 0.16 log points. This dispersion is not an artifact of measurement error: it is nearly absent in settings where uniform pricing is expected, including purchases by different establishments belonging to the same buyer, regulated fuel products, and grocery store purchases. Price gaps are also highly persistent, even during a period of large swings in aggregate inflation. Finally, geographic distance between buyer and seller and the mode of payment (cash or credit) account for negligible fractions of observed price variation, suggesting that cross-buyer price differences primarily reflect markup variation rather than differences in the marginal cost of supplying different buyers.

Our second empirical contribution is to document the sources of this markup hetero-

¹See e.g., Grennan (2013), Davis et al. (2013), Marshall (2020), Fontaine et al. (2020), and Alviarez et al. (2025). We discuss this literature below.

geneity. The dominant source of markup heterogeneity is match-specific: buyer-seller pair identity alone absorbs 43 percent of cross-buyer variation in price gaps. Prices are lower for larger buyers, longer-standing relationships, and higher transaction volumes, but these observable buyer-seller characteristics jointly explain at most 10 percent of cross-buyer variation in price gaps. We also provide evidence that nonlinear price-quantity menus play a limited role in accounting for observed price dispersion: quantity variation is large even among buyer-product pairs with stable prices, and price dispersion is similar among buyers purchasing identical quantities. These findings point to unobserved match-specific factors — such as heterogeneous outside options, switching costs, or relationship-specific bargaining power — as the main source of markup differences across buyers.

Guided by these findings, we study how markup dispersion across buyers of intermediate inputs affects aggregate efficiency. We build on a large literature that links misallocation to markup dispersion (see e.g., [Peters 2020](#), [Baqae and Farhi 2020](#), [Edmond et al. 2023](#), [Pellegrino 2023](#), [Osotimehin and Popov 2023](#)). While the models in this literature are flexible enough to allow for markups to vary across buyer-seller pairs, a standard assumption when mapping them to data is that firms charge common markups across all buyers. We quantify how relaxing this assumption changes the extent of misallocation in these models.

Our model features firms that produce differentiated products using intermediate inputs and a single factor in fixed aggregate supply, and sell their output to other firms and final consumers. We introduce buyer-specific markups as exogenous wedges that match cross-buyer price differences in the data, under the assumption, motivated by our data, that marginal costs do not vary with the identity of the buyer. This allows us to measure markup dispersion across buyers directly from price dispersion, without estimating markup levels for each product-buyer pair, which typically requires strong assumptions (see e.g., [De Loecker et al. 2016](#) and [Dhyne et al. 2022a](#)).² Any model that endogenizes these markups would need to be consistent with our finding that observable product and buyer characteristics explain little of the variation in prices. In line with our evidence on price-quantity menus, we assume that buyers can freely choose quantities at a given price. As in [Hsieh and Klenow \(2009\)](#) and [Baqae and Farhi \(2020\)](#), dispersion in wedges generates differences across buyers in the marginal revenue product of the same input, reducing allocative efficiency.

²Relatedly, the evidence on pricing-to-market in international macroeconomics is based on how relative prices of the same good across destinations respond to exchange-rate changes, rather than first having to estimate markups (see, e.g., the references in [Burstein and Gopinath \(2015\)](#)). Our focus, however, is on differences in markup levels across domestic buyers, rather than changes in markups over time.

We calibrate the model to match the production network, intermediate input shares, and revenue-cost ratios observed in the Chilean data. We consider a range of elasticities of substitution drawn from the literature. Eliminating buyer-level markup dispersion for manufactured intermediate inputs, while holding product-level average markups fixed, increases aggregate productivity by 2.4 to 6.8 percent (normalized by the share of sales with affected markups), depending on the elasticities of substitution. When we also eliminate product-level markup dispersion for manufactured inputs, the normalized productivity gain ranges from 3.8 to 10.6 percent. Buyer-level markup differences thus account for 30 to 64 percent of the total gains from equalizing markups across both buyers and products. Ignoring buyer-level markup dispersion can thus substantially understate the efficiency costs of markup distortions in production networks.

Related literature. Most evidence on price dispersion comes from consumer prices. For example, [Coibion et al. \(2015\)](#), [Kaplan and Menzio \(2015\)](#), [Kaplan et al. \(2019\)](#), and [Sangani \(2025\)](#) study consumer price dispersion within and across retail stores, while [DellaVigna and Gentzkow \(2019\)](#) and [Daruich and Kozlowski \(2023\)](#) document the widespread use of uniform prices within large retail chains. Evidence of cross-buyer price dispersion for intermediate inputs is largely limited to specific industries or firms. [Grennan \(2013\)](#) studies price dispersion across buyers of medical devices, [Davis et al. \(2013\)](#) examine electricity prices across space and time, [Marshall \(2020\)](#) studies wholesale prices paid by different New York restaurants for the same good on the same day, and [Macedoni and Mattana \(2023\)](#) quantify dispersion in firm-to-firm prices for life-saving equipment sold by Viking, a Danish multinational.

Other papers study dispersion of unit values in trade data. [Fontaine et al. \(2020\)](#) use French customs data to document differences in unit values across buyers within CN8 product categories (the tariff lines used by EU customs) sold by the same exporter. As they note, this variation may partly reflect product heterogeneity within CN8 categories. In contrast, we use a much narrower product definition based on product codes and descriptions from firms' billing systems. When we aggregate our products to the HS8 tariff lines used by Chilean customs, price dispersion within product-seller pairs across buyers is about eight times larger than under our narrower product definition. [Alviarez et al. \(2025\)](#) document differences in unit values across US importers purchasing HS10 product categories from the same foreign supplier, while [Ignatenko \(2023\)](#) documents dispersion in unit values across importers over time in Paraguay. These papers study price dispersion through models of oligopoly, oligopsony, or bilateral bargaining. [Fontaine et al. \(2023\)](#) develop a model of Bertrand competition in which buyer-specific markups arise

because buyers are matched with different set of suppliers. Our empirical contribution to this literature is threefold: we characterize price dispersion across domestic buyers for a comprehensive sample of manufactured products; we provide evidence that price gaps primarily reflect markup differences; and we relate these price gaps to observable characteristics of buyers, sellers, and transactions.

Our paper also contributes to a recent literature quantifying the aggregate productivity losses from price dispersion across buyers. Closest to us is [Dhyne et al. \(2022b\)](#), who study the aggregate gains from eliminating markup dispersion across buyers of the same product and infer buyer-specific markups by combining a model of endogenous markups with market-share data. In contrast, we infer differences in markups across buyers directly from product-buyer-specific prices.³ [Bornstein and Peter \(2025\)](#) study misallocation due to markup dispersion across products and buyers, contrasting the implications of linear and non-linear pricing. They provide evidence of non-linear pricing by supermarkets based on price differences across consumer goods with different package size.⁴ Relative to this literature, we focus on prices set by manufacturing firms and treat goods with different descriptions and codes, including package size, as distinct products. Using variation in prices of the same product across buyers, we find that sellers price discriminate across buyers, but that buyers frequently purchase different quantities of a product at the same price. Motivated by this evidence, our model of firm-to-firm production networks assumes linear pricing within product-buyer pairs.⁵

The remainder of the paper is organized as follows. Section 2 introduces the data. Section 3 reports our empirical findings. Section 4 presents our quantitative framework and analytical results. Section 5 reports the results of our counterfactual exercises, and the final section concludes.

2 Data

In this section we describe the dataset and provide basic summary statistics.

³[Koike-Mori and Martner \(2025\)](#) quantify the importance of resource misallocation within the firm for measured productivity growth in Chile by extending the non-parametric approach of [Baqee and Farhi \(2020\)](#) to allow for multi-product firms and joint-production.

⁴Relatedly, [Lorenzini and Martner \(2025\)](#) develop a general equilibrium model of price discrimination by firms along supply chains, motivated by facts on price discrimination by Chilean producers that complement our empirical findings.

⁵The firm-to-firm trade models with endogenous markups in [Dhyne et al. \(2022b\)](#), [Fontaine et al. \(2023\)](#), and [Alvarez et al. \(2025\)](#) also assume that buyers can choose different quantities at a given price. Our model therefore nests these frameworks, even though we do not explicitly model the determination of markups.

2.1 Data description

Our analysis is based on price and quantity data from domestic transactions among firms in Chile. Starting in 2017, the Chilean Internal Tax Service (SII for its acronym in Spanish) mandated the use of electronic invoices for all firm-to-firm sales. Our dataset encompasses all invoices dated from June 2021 to December 2023, and was obtained through the Central Bank of Chile.⁶

The invoices comprise two main sections: the “Heading” and the “Detail”. The Heading section includes essential information regarding the invoice, such as the date, tax identifiers for both the seller and the buyer, the municipalities where they are located, and the payment terms (cash or credit). The Detail section contains specific information about the products sold. Each product is represented by an individual entry within this section. Entries include a description which varies in the level of detail (e.g. “Maule Graphics GC1 315GR 77X110 2500 PL” vs. “Metal plaque”) and a product code (e.g. “ED003C3V59025A”)⁷ which identify the items being traded. Additionally, entries include price, quantity, and any applicable discounts or surcharges for each product.⁸ For 90% of the sales, we have information on the units in which quantities are denominated. Finally, the Central Bank of Chile assigns each firm to a 3-digit ISIC Rev. 4 sector, and each product to broad categories according to the Central Product Classification (CPC).

In what follows, we use the terms “firm” to refer to a unique tax ID, and “establishment” to refer to a tax ID-municipality duplet. We treat all establishments of the same firm as a single seller or buyer. Our set of sellers includes all manufacturing firms, while our set of buyers includes all non-government operated firms in Chile (our set of buyers does not include households).⁹ Because we focus on firm-to-firm sales, we treat each retailer and wholesaler as a single buyer of the manufacturing products. Hence, our price

⁶This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities, by virtue of collaboration agreements signed with these institutions. To secure the privacy of workers and firms, the CBC mandates that the development, extraction and publication of the results should not allow the identification, directly or indirectly, of individuals or firms. Officials of the CBC processed the disaggregated data. All the analysis was implemented by the authors and did not involve nor compromise SII. The information contained in the databases of the SII is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.

⁷This product code corresponds to fully coated folding white boxboards used for the production of medical and healthcare packaging.

⁸For an invoice to be valid, the Chilean Internal Tax Service requires the seller’s and buyer’s tax ID’s, the text description of goods sold, the total amount paid, and a date. More information can be found (in Spanish) at https://www.sii.cl/factura_electronica/formato_dte.pdf.

⁹In our baseline sample of uniquely identified products, described below, 31% of firm-to-firm manufacturing sales are to other manufacturing firms, 54% are to firms Retail and Wholesale, and 15% are to firms in other sectors.

dispersion measures exclude any variation in prices that may arise when retailers and wholesalers resell manufactured products to other firms or final consumers.

We identify unique products using product descriptions and codes. These codes can take different forms, such as EAN (European Article Number) codes or internal codes used by the seller, and are only available to us starting in June 2021.¹⁰ We focus on non-oil manufacturing products that are assigned codes consisting of at least four characters.¹¹ In our baseline analysis, we define a unique product as a product code selling establishment combination. Thus, identical product codes issued by the same seller from different establishments are considered distinct products.¹² For each seller, we discard codes associated with multiple descriptions, as well as descriptions that are associated with multiple codes. We identify unique products using both the descriptions and the codes so as to minimize concerns about descriptions being vague or internal product codes being coarse. We conduct robustness exercises using the smaller sample of products for which EAN product codes are available.

We take the following additional steps to clean the raw data. First, we address potential typos in the invoices by removing transactions that exceed 100 billion pesos (approximately 125 million USD) or that display non-positive prices or quantities. Second, we discard invoices in which the tax ID of the seller matches that of the buyer. Third, we eliminate invoices that were later cancelled or deemed invalid. Fourth, we merge the invoices with the firm's wage and tax statements and discard firms that do not declare any paid workers. Finally, to be conservative in our estimates of price dispersion, we eliminate transactions where the recorded price deviates from the average price recorded in other transactions of the product during the month by a ratio that is above $\exp(1)$ or below $\exp(-1)$.¹³ Appendix B describes these datasets and cleaning steps in detail.

2.2 Summary statistics

Our goal is to analyze price dispersion across buyers who purchase the same product within a narrow period of time. Table 1 reports summary statistics for the sample of

¹⁰The Central Bank of Chile has a data agreement with the Chilean Internal Tax Service, which incorporates several tax forms and documents. Not all the data in the Chilean Internal Tax Service is part of this agreement.

¹¹The median product in our sample consists of 8 characters. Table 4 shows that the extent of price dispersion across buyers is very similar if we only consider products with 8 or more characters.

¹²When units are observed, all transactions of a given product have the same units. Our baseline sample includes products for which we do not observe units across all transactions, but our measures of price dispersion are fairly insensitive to dropping these products.

¹³This last step eliminates transactions accounting for less than 1% of sales in our sample. Our results are largely unchanged if we use a more conservative threshold of $\exp(0.4)$.

invoices issued by manufacturing firms in June 2023.¹⁴ These statistics do not vary much across months between June 2021 and December 2023.

The first column provides an overview of the entire manufacturing sample, consisting of 4.6 million invoices that amount to 5.0 billion USD in sales. The second column focuses on our baseline sample of 372,260 products that we can uniquely identify. This sample covers 2.1 million invoices and 10.7 million transactions, representing almost half of total manufacturing sales (2.3 USD billion). Even though the fraction of sellers using product codes is relatively small (4,299 out of 45,256), they account for almost half the sales and the invoices.

Products purchased by more than one buyer in a month account for 58% of sales of our baseline sample. The average product in our sample has 10.5 buyers, with each buyer making repeat purchases of the same product 2.8 times on average during June 2023. When focusing on products with multiple buyers, the average number of buyers increases to 26.4, though the median of 4 buyers suggests a right-skewed distribution.

The third column of Table 1 reports summary statistics for the sample of products with EAN codes. These products constitute a small portion of the total sales (0.3 billion USD) and are used by only 226 sellers. Products with EAN codes are more frequently sold to multiple buyers and have a higher average number of buyers.

Table 2 shows industrial composition and sectoral inflation rates in our baseline sample and in the official manufacturing PPI data. The largest sector in our sample is ‘Food’, followed by ‘Chemicals’. Importantly, the correlation between sectorial shares in our data and in the PPI is very high (0.97). The level and cross-sectoral variation in inflation rates are also very correlated in our data and in the PPI data.

3 Empirical findings

In this section we report our central empirical findings. We first present our baseline measures of price dispersion and analyze how dispersion varies across time, sectors, products, and sellers. We then show that observed price gaps across buyers are not driven by transportation costs, mode of payment, or by staggered price changes across buyers. Next, we investigate how price gaps vary by characteristics of the buyer and the buyer-seller relation. Finally, we examine how price gaps relate to transaction volumes and assess whether price gaps in our data are consistent with fixed price/quantity menus.

¹⁴We focus on June 2023 in our baseline because it was preceded by a year of relatively low PPI inflation in Chile. Specifically, manufacturing PPI inflation was -3.1% between June 2022 and June 2023.

Table 1: Summary statistics

	All	Baseline	EAN codes
Sales (USD bill.)	5.0	2.3	0.3
Invoices (mill.)	4.6	2.1	0.4
Transactions (mill.)	22.3	10.7	2.7
Sellers	45,256	4,299	226
Buyers	370,585	146,855	53,432
Products		372,260	45,680
Buyer-seller pairs		412,486	65,882
Buyer-product pairs		3,898,999	658,447
Buyers per product		10.5	14.4
Transactions per buyer-product		2.8	4.1
Products with more than one buyer:			
Share of sales		0.58	0.80
Buyers per-product		26.4	25.7
Buyers for median product		4	4
Transactions per buyer-product		5.9	9.7
Transactions for median buyer-product pair		2	3

Notes: The table reports summary statistics for different samples of invoices. 'All' refers to all non-oil manufacturing invoices in June 2023. 'Baseline' refers to our baseline sample, which identifies unique products and also cleans the data as described in the text. 'EAN codes' refers to the sample of invoices where products can be identified using EAN codes.

3.1 Price dispersion across buyers

Let $p_{ib\tau}$ denote the price (after discounts and surcharges) of product i purchased by buyer b in transaction τ .¹⁵ Denote the price gap (in logs) for transaction τ by

$$\mu_{ib\tau} \equiv \log(p_{ib\tau}/p_i), \quad (1)$$

where p_i denotes some average of prices paid by buyers of product i . Our results in this section do not depend on the value of p_i unless when specified.

We start by showing that for the typical product, most of the dispersion in price gaps within a month is across buyers rather than across transactions for the same buyer. To do

¹⁵In our data, 20% of total sales receive a discount, 2% of sales have a surcharge, and 5% of our sales receive both a discount and a surcharge. We focus on final prices after discounts and surcharges as they correspond to the effective transaction prices.

Table 2: Comparison with official PPI data

	Share of sales		Inflation 06/22-06/23	
	Sample	PPI	Sample	PPI
Food	49.3	34.3	6.1	6.4
Beverages	7.7	10.3	11.3	7.2
Textiles	0.6	0.0	6.6	NA
Wood products	4.4	5.0	-2.1	-7.6
Paper	8.5	8.7	2.0	-20.7
Printing	0.1	1.7	11.3	9.7
Chemicals	11.0	11.2	-16.3	-15.7
Pharmaceutical	2.7	3.1	7.0	7.0
Plastic, rubber	3.4	4.3	-5.2	-6.9
Non-metallic minerals	5.2	4.8	11.3	9.3
Metals	3.4	2.6	-21.6	-9.5
Metal products	1.3	7.9	-2.2	-0.3
Electric equipment	0.5	0.6	2.3	7.6
Machinery and equip	0.5	3.6	18.1	6.7
Furniture	1.1	1.9	2.3	5.1
Other	0.2	0.0	16.6	NA
Correlation	97.4		70.9	

Notes: The first panel reports the share of each manufacturing 2-digit ISIC sector in total ‘non-oil’ manufacturing sales. The first column reports the shares for the baseline sample, and the second column reports the shares used by the Chilean Institute of National Statistics for the computation of the Manufacturing Producer Price Index (PPI). The second panel reports the inflation between June 2022 and June 2023 for each 2-digit ISIC sector. The third column reports inflation in our sample, computed as a weighted average of log-price changes with weights based on sales shares in the invoices. The fourth column reports the official PPI inflation in each sector.

so, we regress $\mu_{ib\tau}$ on a full-set of product-buyer fixed effects for the sample of products purchased by multiple buyers during June 2023. Table 3 reports the partial R^2 of these regressions relative to a reduced model that only includes product-level fixed effects (no buyer fixed effects).¹⁶ The partial R^2 ranges between 0.87 and 0.93 depending on the weighting schemes, showing that buyer fixed effects absorb a large portion of the variation in prices across transactions within the month. Results are similar for other months.

Given that there is little variation in prices across transactions within a month for a given product-buyer, we aggregate prices by product-buyer across transactions in each month “ t ” as $p_{ibt} \equiv [\sum_{\tau \in t} p_{ib\tau} q_{ib\tau}] / [\sum_{\tau \in t} q_{ib\tau}]$, and define the price gap for the ib pair as $\mu_{ibt} \equiv \log(p_{ibt}/p_{it})$, where p_{it} is an average of p_{ibt} . Unless noted otherwise, we focus

¹⁶The partial R^2 equals $[1 - SSR^F / SSR^R]$, where SSR^F and SSR^R denote the sum of squared residuals of the full and the reduced model, respectively. This statistic gives the proportion of the variation explained by the full model that cannot be explained by the explanatory variables in the reduced model. Throughout Sections 3.1 to 3.1 and 3.4, the reduced model used for the computations of the partial R^2 is a simple regression of price gaps on product fixed effects.

on $t = \text{June 2023}$ and omit subscript “ t ” to streamline notation. We denote the vector of observed price gaps for different buyers of product i by μ_i .

Table 3: Price dispersion within and across buyers

	(1)	(2)	(3)
Partial R2	0.87	0.89	0.93
Observations	9,843,216	9,843,216	9,843,216
Weights	None	Sales	Quantity

Notes: The Table reports the partial R-squared of the model $\mu_{ib\tau} = FE_{ib} + \epsilon_{ib\tau}$ relative to the reduced model $\mu_{ib\tau} = FE_i + \epsilon_{ib\tau}$, where FE_i and FE_{ib} are a full set of product-level, and product-buyer level fixed effects. The regressions are estimated by OLS on the sample of products that were purchased by more than one buyer in June 2023. In Column 1, each transaction receives an equal weight. In Column 2 transactions are weighted by their sales. In Column 3, each product receives an equal weight, and transactions of the same product are weighted according to the transaction’s share in the product’s quantity.

We calculate three measures of price dispersion for multi-buyer products: (i) the quantity-weighted standard deviation of price gaps across buyers, $\sigma_i \equiv \text{stdev}(\mu_i)$, (ii) the quantity-weighted interdecile range of price gaps across buyers (in log differences), $idr_i \equiv P_{90,W}(\mu_i) - P_{10,W}(\mu_i)$, and (iii) the quantity-weighted interquartile range, $iqr_i \equiv P_{75,W}(\mu_i) - P_{25,W}(\mu_i)$.¹⁷ Under the assumption that marginal costs do not vary by buyer identity, quantity weights correspond to cost weights. We use cost weights because the extent of misallocation arising from markup dispersion depends on cost-weighted, rather than sales-weighted, dispersion (see e.g. [Edmond et al., 2012](#) or [Baqae and Farhi, 2020](#)).

Baseline statistics on price dispersion Table 4 presents distributions of these statistics across products, weighting each product by its sales. The first column considers the interdecile range of price gaps, idr_i . The average (median) interdecile range is 0.16 (0.12) log points, with significant variation across products (the interdecile range for the product in the 75 percentile is 0.24). Columns 2 and 3 display the distributions of interquartile range, iqr_i , and standard deviations, σ_i . The average product has an interquartile range of 0.09 and a standard deviation of 0.08. All three measures show substantial dispersion in prices for the same product across different buyers. Uniform pricing is rare in our sample; less than 10% of total sales involve products with identical prices across all buyers. Appendix Figure A.1 shows the distribution of price gaps pooling all products together. For this figure, price gaps are defined relative to the quantity-weighted average price of each product. We weight the distribution by quantities within products and by sales across

¹⁷In these definitions, $\text{stdev}(\mu_i) = \sqrt{\sum_b \mu_{ib}^2 s_{ib} - [\sum_b \mu_{ib} s_{ib}]^2}$, with $s_{ib} \equiv \frac{q_{ib}}{\sum_b q_{ib}}$, and $P_{x,W}(\mu_i)$ denotes the quantity-weighted x th percentile.

products. The standard deviation of price gaps is 0.10, and the interdecile range is 0.19, both close to our baseline measures in Table 4.

Does using our narrow definition of products—based on detailed invoice codes—make a difference in measuring price dispersion? For a sample of products (representing 92% of sales) we can link the CPC categories assigned by the Central Bank of Chile to the product categories in the 8-digit Harmonized System (HS8) used by Chilean customs. The level of disaggregation of HS8 categories is similar to that used in previous work on price dispersion across buyers based on international trade unit value data.¹⁸ Since CPC categories are typically coarser than HS8 categories, most products in our sample are concorded to multiple HS8 categories. However, we can concord a subset of these products with a unique HS8 category and a common unit of measurement across transactions, representing 11% of sales. For this restricted sample, the median σ_i is 0.44 if we define products as an HS8-selling establishment combination, while the median σ_i is 0.06 if we use our baseline product definition. This underscores the importance of using detailed product codes from invoices to measure price dispersion across buyers of the same product.

Table 4: Within-product dispersion of price gaps

	Baseline			Unweighted Transactions		Long codes	EAN
	idr_i	iqr_i	σ_i	σ_i	σ_i	σ_i	σ_i
P1	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P5	0.00	0.00	0.00	0.00	0.00	0.00	0.01
P10	0.00	0.00	0.01	0.01	0.01	0.00	0.02
P25	0.02	0.00	0.03	0.04	0.03	0.02	0.05
P50	0.12	0.04	0.06	0.07	0.07	0.06	0.08
P75	0.24	0.13	0.11	0.12	0.12	0.11	0.12
P90	0.38	0.23	0.16	0.18	0.17	0.18	0.19
P95	0.49	0.34	0.21	0.24	0.21	0.22	0.23
P99	0.84	0.62	0.33	0.37	0.34	0.32	0.37
Mean	0.16	0.09	0.08	0.09	0.08	0.08	0.10

Notes: idr_i and iqr_i respectively denote the (quantity weighted) interdecile range and interquartile range for product i . σ_i is the quantity-weighted standard deviation of price gaps across buyers of product i . The columns under baseline report the distributions idr_i , iqr_i and σ_i , where the unit of observation is a product i , and products are weighted by their sales and the prices are aggregated at the buyer level and include all June 2023 transactions; $p_{ib} = [\sum_{\tau} p_{ib\tau} q_{ib\tau} / \sum_{\tau} q_{ib\tau}]$ for $\tau \in$ June 2023. Column ‘Unweighted’ uses unweighted standard deviations. Column ‘Transactions’ uses transaction level price gaps $\mu_{ib\tau}$ instead of aggregating at the buyer level. Column ‘Long Codes’ uses the sample of products with codes that are at least eight characters in length. Column ‘EAN’ uses the sample of products that can be identified by an EAN code. The distributions only include products with multiple buyers. P1-P99 represent percentiles of the distribution.

¹⁸For example, [Alvarez et al. \(2025\)](#) uses HS10 categories in the US, [Fontaine et al. \(2020\)](#) use CN8 Categories in France, [Ignatenko \(2023\)](#) uses HS8 categories in Paraguay.

Robustness In Table 4 we re-calculate dispersion of price gaps for different choices of samples and specifications. Column ‘Unweighted’ reports the unweighted distribution instead of the sales-weighted one. Column ‘Transactions’ is based on price gaps by transaction, rather than first aggregating transactions by product-buyer. Notably, these distributions closely resemble the one in our baseline specification. We obtain similar price distributions when we calculate prices at the product-buyer level using transactions aggregated over the full 2023 calendar year (rather than using transactions in June 2023); these additional results are not reported in the paper but are available upon request.

The last two columns consider alternative definitions of unique products. Column ‘Long codes’ uses the subsample of products with codes that are at least eight characters in length, which corresponds to the median product in our baseline. In this sample, the share of products with multiple buyers accounts for 0.52 of sales, and the cross-buyer price dispersion is similar to that in our baseline. Column ‘EAN’ reports results for the smaller sample of products identified by an EAN code. Recall from Table 1 that these products have a greater number of buyers. Price dispersion in this sample is slightly higher than in our baseline sample.

Appendix Table A.1 shows that we also obtain similar values for the interdecile and interquartile ranges in these alternative samples.

Low price dispersion subsamples Could the wide dispersion in prices simply reflect measurement error in our electronic-invoice data? Evidence in Sections 3.3 and 3.4—showing that price gaps persist over time, co-vary across products purchased by the same firm, and strongly with observable buyer–seller characteristics—suggests otherwise. In this section, we further validate the data by showing that dispersion is minimal precisely in the samples and sectors where uniform pricing is expected a priori. Detailed results are relegated to Appendix Table A.2.

We start by showing that different establishments owned by the same buyer firm pay homogeneous prices when purchasing the same good.¹⁹ To do this, we repeat our analysis for the subsample of buyers that have multiple establishments. While the standard deviation of price gaps across buyers in this sub-sample mirrors our baseline results, we find significantly lower dispersion across establishments from the same buyer. For the median product-buyer pair, price dispersion across establishments is zero.

Next, we report price dispersion for fuel and petroleum-derived products which are regulated in Chile and expected to have limited scope for price discrimination across

¹⁹Recall that in our baseline, we treat all establishments of the same firm as a single buyer.

buyers.²⁰ For the median product sold by oil refining firms, the standard deviation of price gaps across buyers is less than 1%. For fuel sold at gas stations and retail locations, the median dispersion across buyers is 1.2%,²¹ which is much lower than the 6.4% in our baseline sample.

Finally, we evaluate price dispersion across buyers of grocery stores. Grocery establishments typically display list prices for their products, with less scope for price discrimination across buyers for the same product.²² For the median product, the standard deviation in prices across buyers is 1.6%, which is far lower than in our baseline sample. Most of this dispersion is over time rather than across buyers. If we focus on sales occurring during the same week, price dispersion across buyers for the median product is zero. In contrast, limiting the period to a week has almost no effect on price dispersion in our baseline sample of manufacturing (firm-to-firm) sales.

Price dispersion and inflation We now evaluate the extent to which price dispersion changes with the overall level of inflation, which displays large swings in our sample period. To do so, we compute the standard deviation of price gaps across buyers, σ_{it} , for every product i and month t between June 2021 and December 2023. Figure 1 plots σ_{it} for the average product in each month, and compares it to the year-to-year change in the Manufacturing PPI. The average standard deviation of price gaps across products changes very little over this period, increasing steadily from 0.07 in 2021 to 0.08 in 2023.²³ In contrast, annual PPI inflation rises from 13.8% in June 2021 to 28.5% in June 2022, and falls to -3.1% in June 2023. Appendix Figure A.2 extends this analysis to the sectoral level, showing that there is no systematic relation between inflation rates and average price dispersion by 2-digit ISIC sectors over time.

Which products have more dispersed prices? We now study how the degree of price dispersion relates to product characteristics. We start by examining variation across manufacturing sectors and across the seller size distribution. Appendix Table A.3 reports the distribution of σ_i for each 2-digit ISIC sector, showing that price dispersion is pervasive across all sectors. The sectors with the most dispersed prices across buyers are ‘Paper’,


²⁰The prices of gasoline, liquified gas, and diesel sold by oil refining companies in Chile are regulated under Law N° 20.765, known as the Fuel Price Stabilization Mechanism.

²¹As in the rest of the paper, this number corresponds to purchases by firms, and does not include households.

²²Retailers price can price discriminate, for example, by offering products of different package sizes, as documented in e.g. Bornstein and Peter (2025).

²³We obtain similar results if we use median rather than average σ_i .

Figure 1: Price dispersion and inflation



Notes: The blue line (left) plots the standard deviation of price gaps, σ_i , for the average product in each month between June 2021 and December 2023. The red line (right) plots the year-to-year change in the producer price index over the same period.

and ‘Chemicals’ (average σ_i of 0.10), while the sectors with least dispersed prices are ‘Metals’ and ‘Wood products’ (average σ_i of 0.05).

To examine the relation between price dispersion and size, we fit regressions of the form

$$\sigma_i = \beta' \mathbf{X}_{s(i)} + FE_{municipality(s)} + FE_{sector(s)} + \epsilon_i. \quad (2)$$

The unit of observation is a product and the dependent variable is the standard deviation of price gaps across buyers of the product. The regressor $\mathbf{X}_{s(i)}$ considers different measures of the size of the seller of the product, while $FE_{municipality(s)}$ and $FE_{sector(s)}$ are fixed effects for seller’s municipality and the seller’s 3-digit ISIC sector.²⁴ All regressions consider the sample of products that are purchased by more than one buyer.

The specification in Column 1 of Table 5 includes only the seller’s municipality and sector, yielding an R^2 of 0.08. Column 2 shows a positive relation between seller size (measured using employment) and price dispersion across buyers. Column 3 uses product sales as a measure of size, while Column 4 uses number of buyers. Finally, Column 5 introduces all these variables simultaneously. While the estimated elasticities are all statistically significantly positive, the R^2 in column 5 is near 0.13. Hence, these product characteristics do not account for a large fraction of the variation in dispersion across products in our sample.

²⁴A municipality is the smallest administrative subdivision in Chile. There are 345 municipalities in Chile, 32 of which are in Santiago. Our data covers 51 3-digit ISIC manufacturing sectors.

Table 5: Price dispersion and product type

	(1)	(2)	(3)	(4)	(5)
$\log \text{emp}_{s(i)}$		0.0078*** (0.0008)			0.0053*** (0.0008)
$\log \text{sales}_{s(i)}$			0.0052*** (0.0004)		0.0016*** (0.0006)
$\log \# \text{buyers}_{s(i)}$				0.0147*** (0.0014)	0.0113*** (0.0016)
Seller's sector FE	Y	Y	Y	Y	Y
Seller's mun. FE	Y	Y	Y	Y	Y
R^2	0.0806	0.1013	0.1042	0.1199	0.1320
Observations	138,200	138,200	138,200	138,200	138,200

Notes: The table reports the results of estimating equation (2) by OLS. $\text{emp}_{s(i)}$ refers to the number of employees working for the seller of product i . $\text{sales}_{s(i)}$ and $\# \text{buyers}_{s(i)}$ respectively denote the value of sales and the number of buyers of product i in June 2023. Seller's sector FE are fixed effects for 51 3-digit ISIC manufacturing sectors covered in our data, and Seller's municipality FE are fixed effects for the 345 municipalities in Chile. Standard errors in parentheses are clustered at the seller level. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level.

3.2 Price gaps and differences in costs of supplying different buyers

Variation in distance between buyers and sellers and in the mode of payment by the buyer (cash or credit) can induce differences in the cost of supplying different buyers, resulting in price dispersion across buyers. In this section we show that these two observable characteristics play a small role in accounting for the observed price gaps in our data.

Shipping costs We start by regressing price gaps on distance:

$$\mu_{ib} = \alpha \mathbb{I}(\log \text{dist}_{ib} = 0) + \beta \log \text{dist}_{ib} + FE_{\text{municipality}(b)} + FE_i + \epsilon_{ib}. \quad (3)$$

Here, dist_{ib} is the kilometer distance between the establishment that sells product i and the buyer,²⁵ $\mathbb{I}(\cdot)$ is an indicator for whether the selling establishment and buyer are in the same municipality, $FE_{\text{municipality}(b)}$ is a fixed effect for the municipality of the buyer, and FE_i is a product fixed effect.

Column 1 in Table 6 shows that the coefficient on the indicator variable is negative:

²⁵We compute the distance between the selling establishment and each buyer as kilometers between the capital cities of their respective municipalities. For buyer-seller pairs from the same municipality, we assign a distance of 1. The distance between the two closest municipalities is 1.3 kilometers. For buyers that purchase the same product from multiple establishments (i.e. unique tax ID's appearing on the invoices purchasing a given product from different municipalities), we use the distance to the average establishment (weighted by purchases of the product). Such product-buyer pairs only comprise 10% of purchases in our sample.

price gaps are on average 0.0078 lower when the seller and the buyer are in the same municipality. The coefficient on distance in Column 2 is positive, indicating that the price is increasing in the distance between seller and buyer.²⁶

While the coefficients on distance are statistically significant, their economic impact is minimal. We illustrate this in several ways. For the average product, the log difference in distance between the nearest and the farthest buyer is 1.3, which, based on the estimates in Column 2, corresponds to only a 0.003 log difference in prices. The largest log distance between two establishments in our sample is 8.4 (4,624 km), which implies a 0.019 higher price relative to a buyer that is just 1 km away from the seller. These price gaps induced by distance are small relative to the overall price gaps reported in Table 4 (recall that for the average product, the interdecile range across buyers is 0.16). Moreover, if we exclude from prices additional surcharges listed in the invoices, the correlation between prices and distance becomes statistically insignificant. This suggests that shipping fees are reflected in the surcharges reported in the invoices, though the surcharges do not account for much of the variation in price gaps in our data.²⁷

Finally, the partial R^2 for distance in regression (3), displayed in columns 1 and 2 of Table 6 is roughly 0.031. In a more flexible specification that includes quadratic terms and allows for 3-digit ISIC sector-specific distance effects the partial R^2 rises to 0.040, indicating that distance does not account for much of the observed dispersion in price gaps across buyers. Price dispersion using prices residualized by distance based on the flexible specification, reported in Column 2 in Table 7, is very similar to that in our baseline.

In sum, shipping costs are not a significant driver of price dispersion across buyers in our data.²⁸

Payment terms To examine the extent to which price gaps reflect differences in payment terms, we leverage information on whether the invoice is paid in cash or credit. This information is available for 64% of transactions, which account for 69% of sales. Among the transactions for which data on payment terms are available, 70% (93% of the sales) are paid in credit.

²⁶Relatedly, Arkolakis et al. (2026) document that transaction volumes decline with the distance between seller and buyer using firm-to-firm sales data from Chile.

²⁷As noted in Section 3.1, our baseline results use prices inclusive of surcharges as these are the effective prices that buyers face. Our other empirical results in Section 3.1 and quantitative results in Section 4 are unchanged if we focus on prices exclusive of surcharges.

²⁸The standard deviation of prices is also very similar in a sample of invoices in which both the seller and the buyer are located within the city of Santiago.

Table 6: Price gaps and differences in costs of supplying different buyers

	(1)	(2)	(3)
$\mathbb{I}(\log \text{dist}_{ib} = 0)$	-0.0078*** (0.0013)	0.0002 (0.0015)	
$\log \text{dist}_{ib}$		0.0023*** (0.0003)	
$\mathbb{I}(\text{credit}_{ib})$			-0.0171*** (0.0009)
Product FE	Y	Y	Y
Buyer mun. FE	Y	Y	N
Partial R^2	0.0311	0.0315	0.0033
Observations	3,629,453	3,629,453	2,354,771

Notes: Columns 1 and 2 report the results of estimating equation (3) by OLS, and standard errors in parentheses are clustered at the buyer-seller level. Column 3 reports the results of estimating equation (4) by OLS, and standard errors in parentheses are clustered at the buyer-seller level. All columns include product-level fixed effects: * significant at the 10% level, **: significant at the 5% level, *** significant at the 1% level.

We consider a regression of price gaps on mode of payment:²⁹

$$\mu_{ib} = \alpha \mathbb{I}(\text{credit}_{ib}) + FE_i + \epsilon_{ib}, \quad (4)$$

where $\mathbb{I}(\text{credit}_{ib})$ is a dummy indicating that buyer b of product i paid in credit. Column 3 of Table 6 shows that the coefficient on the indicator variable is negative: price gaps are on average 0.017 lower when the buyer pays in credit. While this may seem counter-intuitive, it is partly accounted for by the fact that credit is used for bigger transactions and bigger transactions display lower price gaps (we examine this relation in Section 3.4). More importantly, the magnitude of the coefficient is small relative to the observed variation in prices, indicating that the payment terms do not account for much of the observed price gaps. In a more flexible specification that allows for sector-specific effects of payment method, the partial R^2 is only 0.011.³⁰ Finally, Column 3 in Table 7 shows that the dispersion in prices across buyers barely changes when prices are residualized by payment method using this flexible specification.

²⁹We focus on the product-buyer pairs that always exhibit only one payment-term during the month. Buyers that use both cash and credit when purchasing the same product account for only 0.8% of sales.

³⁰The full model in this specification is $\mu_{ib} = \sum_s \alpha_s \mathbb{I}_s(\text{credit}_{ib}) + FE_i + \epsilon_{ib}$, where s is the sector of product i for the sectors used in Table 2.

Table 7: Residualized price dispersion

	Baseline	Residualized distance	Residualized payment
P1	0.00	0.00	0.00
P5	0.00	0.01	0.00
P10	0.01	0.01	0.01
P25	0.03	0.03	0.02
P50	0.06	0.06	0.06
P75	0.11	0.11	0.11
P90	0.16	0.16	0.17
P95	0.21	0.20	0.21
P99	0.33	0.33	0.33
Mean	0.08	0.08	0.08

Notes: The table reports the distribution of the standard deviation of prices across buyers of the same product for different samples. The unit of observation in these distributions is a product i . Column ‘Baseline’ uses our baseline sample as in Table 4. Column ‘Residualized distance’ reports the dispersion in the residuals from a regression like (3) that includes quadratic terms and allows for sector-specific distance effects. Column ‘Residualized payment’ reports the dispersion in the residuals from the regression in residuals from the regression like (4) that includes sector-specific dummies for whether the buyer pays in credit. The distributions only include products with multiple buyers. P1-P99 represent percentiles of the distribution.

3.3 Correlation of price gaps over time

We now evaluate whether price gaps are persistent over time. To do so, we regress the price gap for buyer b , product i and month $t = \text{June 2023}$ on the price gap for the same buyer and product in month $t - k$, controlling for product fixed effects:

$$\mu_{ibt} = \beta_k \mu_{ibt-k} + FE_i + \epsilon_{ibt}. \quad (5)$$

We estimate (5) for lags ranging from $k = 1$ to 12 months. The coefficient β_k measures the sensitivity of the price gap for a buyer-product in period t to the price gap k months before. The left panel of Figure 2 shows that price gaps are very persistent: β_k is above 0.8 for $k = 1$ and 0.6 for $k = 12$. The right panel shows that the partial R^2 is 0.72 for $k = 1$ and 0.43 for $k = 12$. Hence, past levels of price gaps account for most of the variation in current price gaps.³¹


We obtain very similar results in a specification where μ_{ibt} and μ_{ibt-k} are calculated

³¹We also consider an alternative approach to quantify the persistence of price gaps. Specifically, we collapse the data at the monthly level in 2023 and project price gaps on product-buyer fixed effects and product-month fixed effects:

$$\mu_{ibt} = FE_{ib} + FE_{it} + \epsilon_{ibt}.$$

over the subset of product-buyer pairs that exhibit non-zero price changes between t and $t - k$, reported in Figure A.3 of the Appendix. This suggests that the persistence in price gaps is not driven mechanically by nominal price stickiness inherent in long term contracts.

Figure 2: Persistence in price gaps



Notes: The left panel plots the coefficients β^k from equation (5) estimated by OLS, for $k = 1, \dots, 12$. The regression includes product fixed effects. The red lines represent 95% confidence intervals. The right panel plots the corresponding partial R-squared of the regressions. The corresponding regressions results are reported in Appendix Table A.4.

3.4 What drives variation in price gaps?

This section examines how price gaps relate to observable characteristics of buyers and buyer–seller relationships. We first document that price gaps are highly correlated across the different products purchased by a given buyer and relate these gaps to buyer characteristics such as size. We then show that roughly half of the variance in price gaps is captured by buyer–seller pair fixed effects and investigate how these pair-specific gaps vary with observable relationship characteristics. Finally, we show that buyers who purchase larger quantities of a given product pay lower prices, although quantity accounts for only a small share of the overall variation in price gaps.

Price gaps and buyer size Do buyers face correlated price gaps across the various products they purchase? To answer this question, we first regress price gaps (μ_{ib}) on product

The partial R^2 corresponding to the buyer-product fixed effects, FE_{ib} , is as high as 0.8. This indicates that most of the variation in prices gap is across buyers rather than across months for a buyer-product. We obtain similar results for 2022.

fixed effects and buyer fixed effects for the sample of products with multiple buyers. Column 1 of Table 8 reports a partial R^2 of 0.23, indicating that nearly one-quarter of the dispersion in price gaps is explained by buyer identity.

Given the importance of buyer identity, we next investigate which observable characteristics account for these differences. Specifically, we relate price gaps to buyer size, location, and industry by estimating:

$$\mu_{ib} = \beta \log \text{size}_b + FE_{\text{municipality}(b)} + FE_{\text{sector}(b)} + FE_i + \epsilon_{ib}, \quad (6)$$

where $FE_{\text{municipality}(b)}$ and $FE_{\text{sector}(b)}$ are fixed effects for the buyer's municipality and 3-digit ISIC sector.³²

Column 2 of Table 8 includes fixed effects for the buyer's sector and municipality, in addition to product fixed effects, but excludes buyer size. The partial R^2 of these geographic and sectoral controls is 0.040 — about one fifth of the partial R^2 of the buyer fixed effects. Thus, while location and industry matter for prices, they explain a small fraction of the idiosyncratic buyer-level variation.

Columns 3 and 4 introduce two measures of buyer size: total sales and purchases of intermediate inputs (we obtain similar results when using employment as the size measure). Larger buyers pay statistically significantly lower prices, but the implied price gaps are modest. For the average product, moving from the 10th to the 90th percentile of the buyer-size distribution reduces prices by only about 1% (see Appendix Table A.6, Columns 2 and 3). Moreover, the partial R-squares in Columns 3 and 4 of Table 8 are only marginally higher than in Column 2, which includes only sector and municipality fixed effects. Overall, buyer identity (i.e., the buyer fixed effect) explains a sizable share of the variation in price gaps, but observable characteristics —buyer size, sector, and municipality— explain relatively little.

Price gaps and buyer-seller relations Are price gaps correlated across products within buyer-seller pairs, and which characteristics of these relationships account for these gaps? Column 5 of Table 8 reports estimates from a regression of μ_{ib} on product fixed effects and buyer-seller fixed effects, using only buyer-seller pairs that transacted in at least two products. The partial R^2 of this specification is 0.43 — about twice as large as the partial R^2 associated with buyer fixed effects in Column 1. This indicates that price gaps within buyer-seller pairs are highly correlated across products: sellers charge consistently high or low prices across the entire bundle of goods they sell to a given buyer.

³²There are buyers in 158 different 3-digit ISIC sectors.

Table 8: Price gaps, buyer size, and characteristics of the buyer-seller relation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
log sales _b			-0.0019*** (0.0006)				
log purch _b				-0.0021*** (0.0006)			
0 < relation _{s(i)b} < 1 yr						-0.0131*** (0.0011)	
relation _{s(i)b} ≥ 1 yr						-0.0176*** (0.0010)	
log sales _{s(i)b}							-0.0056*** (0.0003)
Product FE	Y	Y	Y	Y	Y	Y	Y
Buyer FE	Y	N	N	N	N	N	N
Buyer X Seller FE	N	N	N	N	Y	N	N
Buyer's mun. FE	N	Y	Y	Y	N	Y	Y
Buyer's sector. FE	N	Y	Y	Y	N	Y	Y
Partial R2	0.225	0.040	0.041	0.041	0.428	0.041	0.047
Observations	3,614,649	3,614,649	3,614,649	3,614,649	3,549,623	3,549,623	3,549,623

Notes: Column 1 reports the results of a regression of log price on product and buyer fixed effects. Columns 2-4 report the results of estimating equation (6) by OLS. sales_b and purch_b respectively denote the sales and intermediate input purchases of buyer *b*. Column 5 reports the results of a regression of log price on product and buyer X seller fixed effects. Columns 6-7 report the results of estimating equation (7) by OLS. “0 < relation_{s(i)b} < 1 yr” and “relation_{s(i)b} ≥ 1 yr” respectively denote indicator variables for whether the first recorded transaction in the data for the corresponding buyer-seller pair occurred over June 2022-May 2023 (i.e. repeated relations as of June 2023, which started over the previous 12 months), and before June 2022. sales_{s(i)b} are the total sales of the seller of product *i* to buyer *b*. Buyer's sector FE are fixed effects for 158 3-digit ISIC sectors covered in our data. Columns 1-4 exclude buyers that only purchase one product, and Columns 5-7 exclude buyer-seller pairs that only transact one product. Standard errors in parentheses are clustered at the buyer level * significant at the 10% level, **: significant at the 5% level, *** significant at the 1% level.

To relate price gaps to observable characteristics of buyer-seller pairs, we estimate regressions of the form

$$\mu_{ib} = \beta' \mathbf{X}_{s(i)b} + FE_{sector(b)} + FE_{municipality(b)} + FE_i + \epsilon_{ib}, \quad (7)$$

where $\mathbf{X}_{s(i)b}$ contains observable characteristics of the relationship between buyer *b* and the seller of product *i*.

We first examine how price gaps vary with the length of the relationship between the seller and the buyer. We measure relationship length as the number of months since their first transaction (starting in June 2021), and construct dummies indicating whether the relationship is new (no previous transactions), started within the past year, or is more than a year old. Column 6 of Table 8 reports a regression of price gaps on these dummies, with

new relationships as the omitted category. For relationships less than a year old (more than a year old), price gaps are 0.013 (0.018) log points lower than for new relationships.³³

Column 7 examines the relation between price gaps and the total volume of sales from the seller of product i to buyer b . Price gaps are lower for buyer-seller pairs with higher sales volumes. However, for the average product, moving from the 10th to the 90th percentile of the buyer-seller sales distribution reduces prices by only 1% (see Appendix Table A.6, Column 5). Consistent with this, the partial R^2 reported in Columns 6 and 7 are very similar to that reported in Column 2, which only includes buyer's sector and buyer's municipality fixed effects.

Finally, Appendix Table A.5 examines whether other observable buyer-seller characteristics help explain price gaps. Buyers that purchase more regularly from a given seller pay slightly lower prices, but this relationship disappears once we control for the size of the buyer-seller relationship. Interestingly, firms that both buy from and sell to each other pay prices that are, on average, 1.4% higher. Prices are also slightly lower for buyers that source a broader range of products from the same seller, although the elasticity is quantitatively small (see Appendix Table A.5). Taken together, these observable buyer-seller characteristics explain little of the variation in price gaps absorbed by buyer-seller fixed effects.

Price gaps and transaction volumes Are price gaps related to the quantity purchased? We pool monthly data from June 2021 to December 2023 and estimate regressions of the form:

$$\mu_{ibt} = \beta \log q_{ibt} + FE_{it} + \epsilon_{ibt}, \quad (8)$$

where q_{ibt} is the quantity of product i purchased by buyer b in month t , and FE_{it} is a product-month fixed effect.

Column 1 of Table 9 shows that buyers purchasing larger quantities of a product face lower price gaps, with an estimated elasticity of -0.015. This implies a 4% lower price when moving from the 10th to the 90th quantity decile for the average product (see Column 6 in Appendix Table A.6). Moreover, the partial R^2 associated with quantity is very low (0.029 in Column 1 of Table 9).

³³The dispersion in price gaps across buyers, as reported in Table 4, is also similar when computed separately for old and new relationships. For each product, we calculate σ_i separately for buyers with relationship lengths above and below the median for that product, and the resulting distributions of σ_i are very similar across the two groups. This suggests that the ability of sellers to price discriminate across buyers does not increase systematically with the age of the relationship.

Columns 2 and 3 of Table 9 add buyer-seller-month and buyer-product fixed effects to examine the relationship between price gaps and quantities across products within buyer-seller pairs, and over time within buyer-product pairs. Larger quantities are still associated with lower prices along both dimensions, but the estimated coefficient on $\log q_{ibt}$ falls to -0.011 with buyer-seller-month fixed effects, and to -0.006 with buyer-product fixed effects. That is, a given buyer receives only modest discounts when purchasing larger quantities. The partial R^2 in these specifications is even lower than in Column 1.

In more flexible specifications that include quadratic terms in $\log q_{ibt}$ and allow for sector-specific quantity coefficients, the partial R^2 associated with quantity remains below 0.04. Thus, although we do find a statistically significant relation between price gaps and quantities, these quantity effects explain only a small share of the residual variation in price gaps in our data.

Table 9: Price gaps and quantities

	(1)	(2)	(3)
$\log \text{quantity}_{ibt}$	-0.0148*** (0.0003)	-0.0111*** (0.0001)	-0.0060*** (0.0001)
Product X time FE	Y	Y	Y
Buyer X seller X time FE	N	Y	N
Buyer X product FE	N	N	Y
Partial R-squared	0.0290	0.0127	0.0044
Observations	97,237,538	97,237,538	97,237,538

Notes: The Table reports the results of estimating equation (8) by OLS. All columns include product-time fixed effects. For the partial R-squared, the reduced model includes all the fixed effects indicated in the corresponding column. Standard errors in parentheses are clustered at the buyer-seller level. We maintain a balanced sample through the three columns by excluding buyer-seller-time triplets where only one product is transacted, and buyer-product pairs that are only observed in one period. * significant at the 10% level, ** significant at the 5% level, *** significant at the 1% level.

Non-linear and sector-specific effects Finally, we estimate a comprehensive specification that jointly incorporates all observable characteristics —buyer size, relationship length and sales, and buyer-product quantities— while allowing for sector-specific linear and quadratic terms. This flexible model, which includes a total of 560 coefficients (excluding product fixed effects), delivers an interdecile range of fitted price gaps of 8% for the average product (see Appendix Table A.6, Column 7). This exceeds the 4% range that we obtain if we fit price gaps using quantities only, indicating that these characteristics are jointly economically meaningful. Nevertheless, the partial R^2 for this specification remains low at 10%, reinforcing the conclusion that observable factors account for only a

limited share of the overall price dispersion.

3.5 Price-quantity menus

In Section 4 we quantify how markup dispersion affects the allocation of inputs across firms. For that analysis, we assume that buyers can purchase any desired quantity of an input at the transacted price; that is, the marginal price is constant over the relevant range. An alternative assumption is that sellers offer nonlinear price-quantity menus to buyers. In this case, prices are not directly allocative other than by making some menus more attractive than others, and observed price gaps need not be informative about misallocation.

Under the strictest form of such menus, there is a one-to-one mapping between quantities and prices. A simple testable implication is that variation in quantities should fully account for variation in prices. The low partial R^2 associated with quantities in the regressions discussed above already suggests that this is not the case. However, changes in prices and quantities may also reflect other shocks that lead producers to adjust the menus they offer.

To test more directly for this form of price-quantity menu, we examine whether across transactions for a given product-buyer pair, a given price is associated with a constant quantity. Specifically, for each product-buyer pair we compute the standard deviation of log prices and log quantities across transactions over a given time window. Column 1 of Table 10 considers transactions in June 2023 and reports the dispersion of quantities for product-buyer pairs whose standard deviation of prices is below 1%. If one-to-one quantity-price menus were the norm, we would expect little dispersion in quantities for pairs that consistently pay the same price. Instead, we observe substantial variation in quantities even among these pairs.

Columns 2 and 3 of Table 10 report log changes in quantities for each product-buyer pair between May and June 2023, and between March and June 2023, restricting attention to pairs whose price change by less than 1% over each of these periods. Once again, we find large changes in quantities for product-buyer pairs with almost no price movement.³⁴

The evidence above suggests that one-to-one price-quantity schedules are not prevalent in our data. This does not rule out the possibility of tiered schedules with flat unit-price regions. We show, however, that there is wide dispersion in prices across buyers purchasing identical quantities, which is difficult to reconcile with price-quantity menus that assign prices primarily as a function of quantity. Column 4 reports the cross-buyer

³⁴We obtain similar results for the December 2022-June 2023 period.

standard deviation of prices for the subset of product-buyer pairs in which the quantity transacted in June 2023 is exactly one unit. This subsample accounts for 3 percent of our baseline sales. Despite the absence of any quantity variation across buyers, price dispersion in this restricted sample is similar to that in our baseline sample; see Table 4. Taken together, these results suggest that fixed price-quantity menus between buyers and sellers are not a central driver of price dispersion in our data. This evidence supports our focus on price discrimination across buyers, rather than on menu-driven nonlinear pricing, as the main source of observed price gaps.

Table 10: Quantity-price menus

	$\text{stdev}_{ib}(\log q_{ibt})$	$ \log q_{ibt} - \log q_{ibt-1} $	$ \log q_{ibt} - \log q_{ibt-3} $	$\text{stdev}_i(\log p_{ib} q_{ib} = 1)$
p1	0.00	0.00	0.00	0.00
p5	0.00	0.01	0.01	0.00
p10	0.01	0.04	0.04	0.00
p25	0.20	0.12	0.16	0.01
p50	0.44	0.30	0.40	0.05
p75	0.63	0.69	0.78	0.09
p90	0.85	1.25	1.43	0.20
p95	1.01	1.79	1.99	0.34
p99	1.54	3.24	3.40	0.57
Mean	0.45	0.52	0.62	0.08

Notes: The first column reports the distribution of the standard deviation of quantities across transaction for a given product-buyer pair, in which the standard deviation of prices is less than 0.01. The distribution excludes product-buyer pairs that record only one transaction. The second and third columns respectively report the distribution of the log-change in quantities for each product-buyer pair between May and June 2023 and between March and June 2023 for pairs that experienced a price change of less than 1% during these periods. The unit of observation in these distributions is a product-buyer pair ib . The distributions are sales weighted. The last column reports the standard deviation of prices across buyers for the subset of product-buyer pairs where the quantity transacted in June 2023 is exactly one unit. The unit of observation in this distribution is a product i .

3.6 Taking stock

The results in Section 3 establish a set of novel facts of how intermediate input prices vary across buyers, which are summarized as follows.

Price dispersion across buyers of the same product is large and pervasive. For the typical product sold to multiple buyers, the within-month interdecile range of prices is 0.16 log points. Price dispersion appears across all manufacturing sectors in our sample, and is unlikely to be an artifact of measurement error: it is nearly absent in settings where

uniform pricing is expected — regulated fuel products, grocery stores, and transactions among establishments within the same parent firm — and remains stable across months with very different rates of aggregate inflation.

The dominant source of this dispersion is match-specific. Price gaps are strongly correlated across all products traded within the same buyer-seller pair, indicating that the commercial relationship between two firms is characterized by systematically high or low prices. This buyer-seller-level component absorbs 43 percent of the variation in price gaps, more than twice the share explained by buyer identity alone. Observable characteristics of buyers, sellers, and transactions — including size, sector, location, relationship length, transaction volume, and quantity — are statistically significant predictors of prices but collectively explain only a small fraction of this variation.

Two additional findings help distinguish markup variation from cost-based explanations. First, geographic distance between buyer and seller explain little of the price variation once invoice surcharges — which directly capture shipping fees — are excluded. Second, buyers who pay on credit face lower prices rather than higher ones, ruling out a simple cash-discount explanation. Together, these results suggest that price gaps across buyers reflect differences in markups rather than differences in the cost of supplying different buyers.

Finally, the evidence in Section 3.5 suggests that prices observed are allocative: buyers can purchase different quantities at a given transacted price, rather than facing nonlinear price-quantity menus. Quantity variation is large even among buyer-product pairs with stable prices, and price dispersion is similar among buyers purchasing identical quantities. Together with the evidence that price gaps primarily reflect markup rather than cost differences, this motivates our quantitative framework, in which buyer-specific markups enter as wedges calibrated to match observed price differences across buyers. We do not model the endogenous determination of these markups. Any such model would have to confront the observation that buyer-seller characteristics, including size and transacted quantity, explain only a small fraction of measured price gaps. We turn to this framework in the next section.

4 Model

In this section we describe the model that we use to quantify the aggregate productivity losses resulting from markup dispersion.

4.1 Setup

We consider a closed economy with J sectors indexed by j , each populated by a discrete number of firms. Firms produce differentiated products, indexed by i , using a bundle of intermediate inputs and services from a factor that is in fixed aggregate supply. Firms sell their products to other firms and to final consumers. Markups are exogenous and can differ across products and buyers.

Technologies of final goods: Final consumption C of the representative household aggregates output of different sectors according to

$$C = \left[\sum_{j \in J} [\tilde{\gamma}^j]^{\frac{1}{\eta}} [C^j]^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}.$$

Sectoral consumption C^j aggregates output of products in sector j according to

$$C^j = \left[\sum_{i \in N_j} \tilde{\gamma}_i^{\frac{1}{\rho}} c_i^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}},$$

where c_i denotes quantity of product i in final output and N_j is the set of sector j products.

Technologies of intermediate products: The production function for producer i is:

$$y_i = z_i \left[[1 - \bar{\alpha}_i]^{\frac{1}{\sigma}} l_i^{\frac{\sigma-1}{\sigma}} + \bar{\alpha}_i^{\frac{1}{\sigma}} m_i^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}.$$

Here l_i denotes the use of factor services by producer i , z_i is a productivity shifter, and m_i is a bundle of intermediate inputs:

$$m_i = \left[\sum_{j \in J} [\bar{\omega}_i^j]^{\frac{1}{\eta}} [m_i^j]^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where the sectoral j bundle of intermediate inputs used by producer i is

$$m_i^j = \left[\sum_{i' \in N_j} \bar{\omega}_{i'i}^{\frac{1}{\rho}} x_{i'i}^{\frac{\rho-1}{\rho}} \right]^{\frac{\rho}{\rho-1}},$$

and $x_{i'i}$ denotes the quantity of product i' used in the production of i .

Prices and wedges: Firms minimize costs and set prices as a markup over marginal cost, producing the quantity that is demanded at that price. Let $p_{ii'}$ and p_{ic} respectively denote the price of product i when sold to i' and when sold to final consumers. These prices are given by:

$$p_{ii'} = \mu_{ii'} mc_i; \quad p_{ic} = \mu_{ic} mc_i, \quad (9)$$

where μ_{ii} and μ_{ic} are exogenous markups and mc_i is the marginal cost of production of product i . Let the factor price be the numeraire.

Market clearing: Goods market clearing implies:

$$y_i = c_i + \sum_{i'} x_{ii'} \quad \forall i, \quad (10)$$

and factor market clearing implies:

$$\sum_i l_i = \bar{L}. \quad (11)$$

Equilibrium: Given technologies and wedges $\{\mu_{ii}, \mu_{ic}\}$, an equilibrium is a set of goods prices $\{p_{ii'}, p_{ic}\}$, intermediate input quantities $\{x_{ii'}\}$, factor quantities $\{l_i\}$, outputs $\{y_i\}$, and final output quantities $\{c_i\}$, such that (i) firms minimize costs and set prices to $p_{ii'} = \mu_{ii'} mc_i$ and $p_{ic} = \mu_{ic} mc_i$; (ii) final producers minimize costs; and (iii) all markets clear. Appendix C characterizes the equilibrium.

4.2 Solving for changes in allocations

We first provide a system of equations to evaluate how equilibrium prices and quantities change in response to exogenous changes in markups. Denote the ratio of a variable in the new relative to the initial equilibrium by $\hat{X} \equiv \frac{X'}{X}$. Changes in product level prices can be expressed as:

$$\hat{p}_{ii'} = \hat{\mu}_{ii'} \widehat{mc}_i; \quad \hat{p}_{ic} = \hat{\mu}_{ic} \widehat{mc}_i, \quad (12)$$

with

$$\widehat{mc}_i = \left[1 - \alpha_i + \alpha_i \hat{q}_i^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (13)$$

Here, $\alpha_i \equiv \frac{q_i m_i}{l_i + q_i m_i}$ is the cost share of materials for product i in the initial equilibrium, and $\hat{q}_i$ is the change in the price of the input bundle used in the production of i :

$$\hat{q}_i = \left[\sum_j \omega_i^j [\hat{q}_i^j]^{1-\eta} \right]^{\frac{1}{1-\eta}}, \quad \text{where} \quad \hat{q}_i^j = \left[\frac{1}{\omega_i^j} \sum_{i' \in j} \omega_{i'i} \hat{p}_{i'i}^{1-\rho} \right]^{\frac{1}{1-\rho}}, \quad (14)$$

with $\omega_{i'i} \equiv \frac{p_{i'i} x_{i'i}}{q_i m_i}$ and $\omega_i^j \equiv \sum_{i' \in j} \omega_{i'i}$. Changes in final good prices are given by:

$$\hat{P} = \left[\sum_j s^{jc} [\hat{P}^j]^{1-\eta} \right]^{\frac{1}{1-\eta}}, \quad \text{where} \quad \hat{P}^j = \left[\frac{1}{s^{jc}} \sum_{i \in j} s_i^c \hat{p}_{ic}^{1-\rho} \right]^{\frac{1}{1-\rho}}, \quad (15)$$

with $s_i^c \equiv \frac{p_{ic} c_i}{\sum_i p_{ic} c_i}$ and $s^{jc} \equiv \sum_{i \in j} s_i^c$. Finally, goods market clearing implies:

$$\hat{y}_i = \frac{c_i}{y_i} \hat{p}_{ic}^{-\rho} [\hat{P}^j]^{\rho-\eta} \hat{P}^\eta \hat{C} + \left[1 - \frac{c_i}{y_i} \right] \sum_{i'} s_{ii'} \hat{p}_{ii'}^{-\rho} [\hat{q}_{i'}^j]^{\rho-\eta} \hat{q}_{i'}^{\eta-\sigma} \widehat{mc}_{i'}^\sigma \hat{y}_{i'}, \quad \forall i \quad (16)$$

where $s_{ii'} \equiv x_{ii'} / \sum_{i'} x_{ii'}$. Factor market clearing implies

$$1 = \sum_i s_i^l \widehat{mc}_i^\sigma \hat{y}_i. \quad (17)$$

with $s_i^l \equiv l_i / L$.

The system of equations listed above can be used to solve for changes in aggregate consumption $\hat{C}$ in response to changes in markups. In our counterfactuals, in which factor quantities are fixed, changes in aggregate consumption coincide with change in aggregate productivity. We describe how we assign values to the parameters in Section 5.

Decomposing the effects of eliminating markup dispersion In our counterfactuals, we change *buyer-product* specific markups for a subset of products from their observed level $\mu_{ii'}$ to some uniform level $\bar{\mu}$. Changes in markups are $\hat{\mu}_{ii'}^{bp} \equiv \bar{\mu} / \mu_{ii'}$ and the resulting change in aggregate productivity is $\hat{C}^{bp}$. To isolate the role of markup dispersion across buyers, we first eliminate dispersion in markups across buyers, while keeping the average markup of each product unchanged. Changes in markups in this first step are $\hat{\mu}_{ii'}^b \equiv$

$\bar{\mu}_i / \mu_{ii'}$, where $\bar{\mu}_i \equiv \sum_{i'} \mu_{ii'} s_{ii'}$ is the quantity-weighted average markup of product i across buyers in the initial equilibrium, and the resulting change in aggregate productivity is $\hat{C}^b$. Second, we eliminate dispersion in markups across products, abstracting from markup variation across buyers of the same product (which is the standard assumption in the literature). Changes in markups in this second step are $\hat{\mu}_i^p \equiv \bar{\mu} / \bar{\mu}_i$, and the resulting change in aggregate consumption is $\hat{C}^p$. The overall change in aggregate productivity, $\hat{C}^{bp}$, can be expressed as the product of two terms:

$$\hat{C}^{bp} = \hat{C}^b \times \hat{C}^p. \quad (18)$$

The first component captures the impact of eliminating markup dispersion across buyers, and the second component captures the impact of eliminating markups across products. We quantify the contribution of the first component in the overall productivity change. Before turning to the quantification, we illustrate using a simplified version of our model how reducing dispersion of markups across buyers can increase aggregate productivity.

4.3 Analytic results under simplified network

Consider a simplified network with two sectors labeled u (upstream) and d (downstream). Goods in sector u are produced using factor services and are sold to the d sector. That is, $\alpha_i = 0$ and $c_i / y_i = 0$ for $i \in N_u$. Conversely, goods in sector d are produced using intermediate inputs and sold to final consumers: $\alpha_i = 1$ and $c_i / y_i = 1$ for $d \in N_d$.

To eliminate markup dispersion while keeping the average markup unchanged for intermediate goods $i \in N_u$, we set $\hat{\mu}_{ii'}^b = \bar{\mu}_i / \mu_{ii'}$. For downstream goods $i \in N_d$, we leave markups unchanged, $\hat{\mu}_{ic} = 1$. In Appendix C we show that the change in aggregate productivity is

$$\hat{C}^b = \left[\sum_{i \in N_u} s_i^x \frac{\bar{\mu}_u}{\bar{\mu}_i} \sum_{i' \in N_d} \left[\hat{\mu}_{ii'}^b \right]^{-\rho} s_{ii'} \right]^{-1} \left[\sum_{i \in N_u} s_i^x \sum_{i' \in N_d} \frac{\bar{\mu}_{i'}}{\bar{\mu}_d} \left[\hat{\mu}_{ii'}^b \right]^{-\rho} s_{ii'} \right]^{\frac{\rho}{\rho-1}}, \quad (19)$$

where $\bar{\mu}_u$ and $\bar{\mu}_d$ are the average markups in sectors u and d , and $s_i^x \equiv \frac{\sum_{i'} p_{ii'} x_{ii'}}{\sum_{i \in N_u} \sum_{i'} p_{ii'} x_{ii'}}$ is the share of input i in total input sales. If across downstream firms, markups on final sales are uncorrelated with markups on their purchases ($\bar{\mu}_{i'}$ uncorrelated with $\hat{\mu}_{ii'}^b$ across $i' \in N_d$), and across upstream, firms average markups are uncorrelated with markup dispersion over different buyers ($\bar{\mu}_i$ uncorrelated with $\sum_{i' \in N_d} \left[\hat{\mu}_{ii'}^b \right]^{-\rho} s_{ii'}$ across $i \in N_u$), equation (19)

simplifies to

$$\hat{C}^b = \left[\sum_{i \in N_u} s_i^x \tilde{\zeta}_i \right]^{\frac{1}{\rho-1}}, \quad \text{where} \quad \tilde{\zeta}_i \equiv \frac{\sum_{i'} s_{ii'} \mu_{ii'}^\rho}{[\sum_{i'} s_{ii'} \mu_{ii'}]^\rho} \geq 1 \quad \text{if } \rho > 1. \quad (20)$$

In this case, aggregate productivity rises if markups are initially dispersed across buyers and $\rho > 1$. The rise in C is increasing in the initial markup dispersion and in the magnitude of the elasticity ρ .

Without imposing these strong assumptions, equation (19) shows that $\hat{C}^b$ depends also on the covariance between changes in wedges $\hat{\mu}_{ii'}^b$ and product-level markups $\bar{\mu}_{i'}$. Aggregate productivity may decline if, for example, downstream firms that charge high markups pay low markups on their inputs ($\bar{\mu}_{i'}$ positively correlated with $\hat{\mu}_{ii'}^b$ across $i' \in N_d$). In this case, eliminating intermediate input price dispersion raises input prices and shifts production away from these high-markup downstream firms. In our quantitative exercise, we target $\bar{\mu}_{i'}$ and $\hat{\mu}_{ii'}^b$, which in turn fixes these covariances.

5 Quantitative results

In this section we describe how we map the model to the Chilean electronic invoice data used above and to income statement forms collected by the Chilean tax authority (Form 22 or IDF). These forms provide firm-level data on input expenditures, total costs, revenues, and value added. Appendix B describes these additional datasets. We then quantify the change in aggregate productivity from equating markups across buyers, and across buyers and products.

5.1 Mapping the model to data

Sample of products In our quantification, the network of firms includes the sellers of the products that can be identified uniquely as described in Section 2, as well as the buyers of these products from the manufacturing, retail, and wholesale sectors. As summarized in Appendix Table A.7, in June 2023 there are 237,468 products and 128,748 firms in this network. This results in more than 30 billion possible product-buyer pairs, of which 3,249,102 have positive sales. Our counterfactual exercises thus require solving a large non-linear system of equations. In Appendix E we describe our solution algorithm.³⁵

³⁵Similar to Huneeus (2020), our algorithm manages the system's large dimensionality by exploiting sparsity of firm-to-firm linkages. Instead of utilizing a sparse matrix of all possible product-buyer pairs, the algorithm focuses on buyer-product combinations with positive sales.

Obtaining counterfactual changes in markups The change in markups needed for our counterfactuals is:

$$\hat{\mu}_{ii'}^{bp} = \hat{\mu}_{ii'}^b \times \hat{\mu}_{i'}^p. \quad (21)$$

To calculate the first term, $\hat{\mu}_{ii'}^b \equiv \bar{\mu}_i / \mu_{ii'}$, we assume that marginal costs of production are common across buyers. Substituting $\bar{\mu}_i = \sum_b \mu_{ib} s_{ib}$ we obtain

$$\hat{\mu}_{ii'}^b = \frac{\sum_b \mu_{ib} s_{ib}}{\mu_{ii'}} = \frac{\sum_b \mu_{ib} mc_i s_{ib}}{\mu_{ii'} mc_i} = \frac{\sum_b p_{ib} s_{ib}}{p_{ii'}}. \quad (22)$$

We use our price data presented in Section 2 to compute $\hat{\mu}_{ii'}^b$ in (22), without having to estimate average product markups $\bar{\mu}_i$.

Note that the data on prices p_{ib} required for implementing equation (22) are only available for the set of manufactured products with unique codes. In our counterfactual exercises, we only change markups for products with unique codes, keeping markups for products without codes constant. We also keep markups for all final sales constant, $\hat{\mu}_{ic} = 1$. Products with unique codes account for 31% of total sales (including manufacturing, non-manufacturing, and final sales). Firm-to-firm sales of these products account for 12% of total sales, and the firm-to-firm sales of the subset of these products that have multiple buyers account for 7.1% of total sales.

To calculate the second term, $\hat{\mu}_{i'}^p \equiv \bar{\mu} / \bar{\mu}_i$, we assume that product level markups $\bar{\mu}_i$ are equal for all products sold by the same firm, and set $\bar{\mu}_i$ to match the ratio of revenues to costs for firm that sells product i obtained from income statement forms. We set the counterfactual uniform markup $\bar{\mu}$ as the harmonic average markup on intermediate sales across product with codes, calculated as

$$\bar{\mu} = \left[\sum_{i \text{ w/ code}} \bar{\mu}_i^{-1} \frac{s_i^x}{\sum_{i \text{ w/ code}} s_i^x} \right]^{-1}.$$

Here, s_i^x is the share of product i in total intermediate sales derived from the invoices below, and the summations are across products with codes. This yields $\bar{\mu} = 1.074$, which is similar to the ratio of total revenues to total costs in the economy including products without codes. Note that, since we only change markups for a subset of firms, our counterfactual exercises are not designed to quantify the productivity gains of moving to the first best (which requires setting all markups to one).

Parameterization To evaluate the system in Section 4.2, we need to assign values to the cost share of inputs, α_i , the share of each input i' in input expenditures, $\omega_{i'i}$, the share of

the fixed factors used for the production of each product, s_i^l , the share of each product in final sales, s_i^c , and the fraction of the product that is sold to final consumers and to each other intermediate product i' , $\frac{c_i}{y_i}$ and $s_{ii'}$. We also need to assign values to the elasticities of substitution between factors and materials σ , and to the elasticities of substitution across products and sectors, ρ and η .

Rather than making distributional assumptions to calibrate the shares, we map each product in the model to a product in the Chilean data. A challenge is that, while most firms in the data sell multiple products, data on input use and on sales to final consumers are only available at the firm level. Furthermore, the invoice data identify the buyer of each input, but not how buyers apportion inputs across their multiple products.³⁶ To apportion firm-level inputs to products, we assume that each firm uses the same input bundle and sets the same average markup across its multiple products. We also assume the ratio of sales to final consumers to total sales is the same across all products of the same firm. Appendix D shows that under these assumptions, we can apportion firms' inputs across products in proportion to intermediate sales.

We treat purchases of inputs that are not in the model's network as purchases of factor services.³⁷ We denote by θ_f the share of input purchases of firm f that are included in the model's network by θ_f , by $f(i)$ the identity of the firm that produces product i , and by $\kappa_i \equiv [\sum_b p_{ib} x_{ib}] / [\sum_{i \in f(i)} \sum_b p_{ib} x_{ib}]$ the share of product i sales in total sales of firm $f(i)$. We set $\omega_{i'i}$ as the ratio of purchases of product i' relative to all input purchases by firm $f(i)$ from the invoices. We set $s_{ii'}$ as the product of the quantity share of product i that goes to firm $f(i')$, $\frac{x_{if(i')}}{\sum_{f'} x_{if'}}$, and the share of product i' in $f(i')$ sales, $s_{ii'} = \frac{x_{if(i')}}{\sum_{f'} x_{if'}} \times \kappa_i$.

We obtain the remaining shares after first calibrating the following auxiliary variables. First, from the invoices we compute the share of each product i in total intermediate input sales, s_i^x , and the share of each product in total intermediate input purchases, s_i^m .³⁸ Second, we use the income statement forms to compute the ratio of value added to revenues for each product i , which we denote by v_i . We then obtain the remaining shares as $\alpha_i = \bar{\mu}_i \times [1 - v_i]$, $s_i^c = [s_i^m \frac{1}{1-v_i} - s_i^x] \frac{1-v}{v}$, $s_i^l = [\frac{1-\alpha_i}{\alpha_i} / \frac{1-\alpha}{\alpha}] s_i^m$, $s_i^r = v s_i^c + [1 - v] s_i^x$ and $\frac{c_i}{y_i} = [1 + \frac{1-v}{v} \frac{s_i^x}{s_i^c}]^{-1}$. Here, v and α denote the economy-wide ratio of value added to revenues and the economy-wide ratio of input purchases to total cost. Using the

³⁶That is, the invoices have data on prices and quantities at the product-buyer level p_{ib} and x_{ib} , where the buyer b is identified by its tax ID.

³⁷As noted above, we focus on the network of products for which we observe product identifiers, and additionally include all the buyers of these products in both the Manufacturing and Retail and Wholesale sectors. This implies that, for example, purchases of agricultural inputs are not considered in the network and are lumped with into factor services.

³⁸Specifically, we set $s_i^x \equiv [\sum_{i'} p_{ii'} x_{ii'}] / [\sum_i \sum_{i'} p_{ii'} x_{ii'}]$, $s_i^m = [\sum_{i'} p_{if(i)} x_{i'f(i)}] / [\sum_{f'} \sum_{i'} p_{i'f'} x_{i'f'}] \times \kappa_i$.

previously targeted shares, we obtain these ratios as $1 - v = \left[\sum_i \frac{sm_i}{1-v_i} \right]^{-1} = 0.30$ and $\alpha = \left[\sum_i s_i^m / \alpha_i \right]^{-1} = 0.32$. Table 11 summarizes the calibration of initial shares, and Appendix D provides additional details.

Finally, we set the the elasticity of substitution between factor services and the input bundle to $\sigma = 1$, and assume for our baseline that there is only one sector, $\rho = \eta$. We report our findings for low and high values of ρ : $\rho = 3$ or $\rho = 6$.³⁹ We provide robustness of our quantitative results to these assumptions.

Table 11: Calibration of initial shares

Targeted ratios	
$\bar{\mu}_i$	Ratio of total revenues to total costs of firm $f(i)$, EI.
$\omega_{i'}$	Share of intermediate input i' in input purchases of firm $f(i)$, EI.
s_i^x	Share of input i in total intermediate input sales, EI.
κ_i	Share of product i in firm's $f(i)$ sales, EI.
$s_{ii'}$	Quantity share of intermediate of i going to $f(i')$ multiplied by κ_i , EI.
s_i^m	Share of $f(i)$ in total intermediate input purchases multiplied by κ_i , EI.
v_i	Ratio value added to revenues of firm $f(i)$. IDF, EI.
Non targeted ratios	
α_i	Share of intermediate inputs in total cost: $\alpha_i = \bar{\mu}_i \times [1 - v_i]$.
s_i^c	Share of product i sales in total final sales: $s_i^c = \left[s_i^m \frac{1}{1-v_i} - s_i^x \right] \frac{1-v}{v}$.
s_i^l	Share of the fixed factor that is used for the production of i : $s_i^l = \left[\frac{1-\alpha_i}{\alpha_i} / \frac{1-\alpha}{\alpha} \right] s_i^m$.
s_i^r	Share of product i sales in total sales: $s_i^r = v s_i^c + [1 - v] s_i^x$.
$\frac{c_i}{y_i}$	Fraction of good i that is sold to final consumers: $\frac{c_i}{y_i} = \left[1 + \frac{1-v}{v} \frac{s_i^x}{s_i^c} \right]^{-1}$.
α	Economy wide share of inputs in production cost: $\alpha \equiv \left[\sum_i s_i^m / \alpha_i \right]^{-1} = 0.32$.
v	Economy wide ratio of value added to revenues: $v \equiv 1 - \left[\sum_i s_i^m / [1 - v_i] \right]^{-1} = 0.70$.

Notes: This table describes how we assign values to the model's initial shares under our baseline calibration. $f(i)$ refers to the firm (tax id) that sells product i . 'EI' and 'IDF' indicate that calculations are based on data sourced from the Electronic Invoices and from the Income Declaration Forms respectively.

³⁹The cost of markup dispersion increases with the value of the elasticity of substitution. Hsieh and Klenow (2009) set this elasticity to 3, Baqaee and Farhi (2020) consider values of 4 and 8, while Edmond et al. (2023) consider elasticities between 5.7 and 29.1.

5.2 Productivity gains from eliminating markup dispersion

Table 12 presents results from our counterfactual exercises. In the first panel, we consider changes in markups for firm-to-firm sales of coded products with multiple buyers, which account for 7.1% of total sales (Column 1). In Row 1, we equalize markups across buyers while keeping average product markups unchanged. The change in markups corresponds to the term $\hat{\mu}_{ii'}^b$ in the decomposition (21). Setting the elasticity of substitution $\rho = 3$, the change in aggregate productivity is $\hat{C}^b - 1 = 0.17\%$ (Column 2). The change in aggregate productivity normalized by the fraction of sales with changed markups is $0.17/0.071 \approx 2.4\%$ (Column 3). Setting $\rho = 6$, the normalized change in aggregate productivity is 6.8% (Column 5).

In the second row, we equalize markups both across buyers and across products, corresponding to the term $\hat{\mu}_{ii'}^{bp}$ in (21). The increase in aggregate productivity, $\hat{C}^{bp} - 1$, normalized by the sales share of affected products is 3.8% with $\rho = 3$, and 10.6% with $\rho = 6$. Comparing Rows 1 and 2, the ratio of efficiency gains from eliminating markup dispersion across buyers relative to the total gains from equating markups across buyers and products, $(\hat{C}^b - 1)/(\hat{C}^{bp} - 1)$, is roughly 60% for both values of ρ .

In the second panel of Table 12, we change markups for all product with codes, including those with only one buyer. These products account for 11.9% of total sales. Equalizing markups across buyers (Row 3), the change in aggregate productivity $\hat{C}^b$ is the same as in Row 1, since markup dispersion across buyers is zero for single buyer-products. However, since the sales share of affected products is larger when we include single-buyer products, the normalized change in productivity drops (relative to Row 1) to 1.4% with $\rho = 3$, and to 4.0% with $\rho = 6$. Eliminating markup dispersion across both buyers and products (Row 4), the normalized increase in aggregate productivity is 3.8% with $\rho = 3$, and 13.2% with $\rho = 6$, which is similar to that when we only shock products with multiple buyers. Comparing Rows 3 and 4, the aggregate efficiency gains from eliminating markup dispersion across buyers is about 30% of the total gains from equating markups across buyers and products.

We conclude from these exercises that aggregate productivity gains from eliminating markup dispersion across buyers range between 30% and 60% of the total gains from eliminating dispersion across buyers and products, depending on whether we consider in our counterfactuals changes in markups for single buyer products. This conclusion is robust to plausible values of the elasticity of substitution across products.

Robustness Table 13 reports the results of our counterfactuals under alternative samples and values of the key elasticities. For brevity, we only consider the counterfactual

Table 12: Productivity gains from eliminating markup dispersion

	shocked sales / total sales	Low elasticity: $\rho = 3$ productivity (%) (2)/(1)		High elasticity: $\rho = 6$ productivity (%) (4)/(1)	
	(1)	(2)	(3)	(4)	(5)
Multibuyer products with codes					
1) Equalizing markups across buyers, $\hat{\mu}^b$	0.071	0.17	2.38	0.48	6.80
2) Equalizing markups across buyers and products, $\hat{\mu}^{bp}$	0.071	0.27	3.76	0.75	10.57
Contribution of equating markups across buyers (1 / 2)		0.63		0.64	
All products with codes					
3) Equalizing markups across buyers, $\hat{\mu}^b$	0.119	0.17	1.41	0.48	4.03
4) Equalizing markups across buyers and products, $\hat{\mu}^{bp}$	0.119	0.45	3.79	1.58	13.23
Contribution of equating markups across buyers (3 / 4)		0.37		0.30	

Notes: Rows 1 and 3 report results from the counterfactual that eliminates dispersion in markups across buyers of a given input, while keeping average product markups unchanged. Rows 2 and 4 report results from the counterfactual that eliminates dispersion in markups across buyers and products. The first panel eliminates dispersion for intermediates sales for the set of products that are assigned codes and that have more than one buyer. The second panel eliminates dispersion for intermediates sales for the set of products that are assigned codes.

where we equalize markups across buyers within products for multi-buyer products. The second row reports results under the EAN definition of products. Since only a small fraction of products use EAN codes, the fraction of shocked sales in this exercise drops to 2%. The normalized change in aggregate productivity in this exercise is larger than in the baseline: 3.1% for $\rho = 3$ and 8.5% for $\rho = 6$. This reflects that EAN products, on average, have a higher number of buyers and more dispersed prices across these buyers compared to the baseline products.

Row 3 reports results using our baseline definition of products but restricting the sample to sellers and buyers located within the city of Santiago. As discussed in Section 2, the assumption that the marginal cost of production is independent of the buyer's identity is more likely to hold when the seller and the buyer are within the same city, as this minimizes potential differences in delivery costs across buyers. The normalized changes in aggregate productivity are similar to those in our baseline sample when we set $\rho = 3$ (2.40% vs. 2.38%), and when we set $\rho = 6$ (6.7% vs. 6.8%).

The last two rows show results for alternative values of the elasticities in the production functions, σ and η . In Row 4, we split products with codes and products without codes into two sectors (products with codes are only in the manufacturing sector, while products without codes are mostly in the retail and wholesale sector)⁴⁰, and set the elasticity of substitution between sectors to $\eta = 1$. Lowering this elasticity slightly reduces the normalized change in aggregate productivity (2.14% vs. 2.38% when $\rho = 3$). In Row 5 we lower the elasticity of substitution between factors and materials to $\sigma = 0.5$. This parameterization results in a slightly smaller change in aggregate productivity compared to the baseline (2.28% vs. 2.38% when $\rho = 3$). Overall, our baseline results are not very sensitive to plausible changes in these two elasticities.

6 Conclusion

Price dispersion across buyers is pervasive in manufactured intermediate input markets. Using comprehensive electronic invoice data from Chilean manufacturing firms, we document that, among products sold to multiple buyers, the typical product exhibits an interdecile price range of 0.16 log points across buyers. This dispersion in prices is present across manufacturing sectors, stable over time, and primarily reflects differences in markups rather than differences in cost of supplying different buyers. The dominant

⁴⁰83% of the sales products of without codes are in the retail and wholesale sector.

Table 13: Robustness to alternative samples and calibrations

	shocked sales / total sales	Low elasticity: $\rho = 3$ productivity (2)/(1) (%)		High elasticity: $\rho = 6$ productivity (4)/(1) (%)	
	(1)	(2)	(3)	(4)	(5)
Baseline	0.071	0.17	2.38	0.48	6.80
EAN	0.024	0.08	3.12	0.21	8.46
Santiago	0.071	0.17	2.40	0.48	6.74
$\eta = 1$	0.071	0.15	2.14	0.38	5.39
$\sigma = 0.5$	0.071	0.16	2.28	0.46	6.52

Notes: This table reports the results of our counterfactual that eliminates the dispersion in markups across buyers of a given input under alternative calibrations and samples. ‘Baseline’ refers to our baseline calibration. ‘EAN’ refers to a network that comprises the EAN products all the buyers of EAN products that are in the Manufacturing and the Retail and Wholesale sectors. ‘Santiago’ uses our baseline calibration but limits the network to products and firms located in Santiago. ‘ $\eta = 1$ ’ divides our sample of products into two sectors and sets $\eta = 1$. The last row follows our baseline calibration but sets $\sigma = 0.5$.

source of this dispersion is match-specific: buyer-seller pair identity absorbs about 45 percent of the variation in price gaps, with observable characteristics of buyers, sellers, and transactions explaining up to 10 percent. This match-specific heterogeneity most plausibly reflects differences in outside options, switching costs, or bargaining power across buyer-seller relationships.

Using a workhorse production-network model with buyer-product markup wedges calibrated to match observed price differences across buyers, we show that markup dispersion across buyers accounts for 30 to 65 percent of the aggregate productivity gains from equalizing markups across both buyers and products in manufacturing firm-to-firm input sales. An important area for future research is to develop models with endogenous markups and endogenous production-network links that are consistent with the facts documented here, and to quantify their implications for misallocation.

References

Alviarez, Vanessa I, Michele Fioretti, Ken Kikkawa, and Monica Morlacco, “Two-Sided

- Market Power in Firm-to-Firm Trade,” Working Paper 31253, National Bureau of Economic Research July 2025.
- Arkolakis, Costas, Federico Huneeus, and Yuhei Miyauchi**, “Production Network Formation, Trade, and Welfare,” NBER Working Papers 30954, National Bureau of Economic Research, Inc January 2026.
- Baqae, David Rezza and Emmanuel Farhi**, “Productivity and Misallocation in General Equilibrium,” *The Quarterly Journal of Economics*, 2020, 135 (1), 105–163.
- Bornstein, Gideon and Alessandra Peter**, “Nonlinear Pricing and Misallocation,” *American Economic Review*, November 2025, 115 (11), 3853–3908.
- Burstein, Ariel T. and Gita Gopinath**, “International Prices and Exchange Rates,” in Kenneth Rogoff Elhanan Helpman and Gita Gopinath, eds., *Handbook of International Economics*, Vol. 4, Elsevier, 2015, chapter 7, pp. 391 – 451.
- Coibion, Olivier, Yuriy Gorodnichenko, and Gee Hee Hong**, “The Cyclicalities of Sales, Regular and Effective Prices: Business Cycle and Policy Implications,” *American Economic Review*, March 2015, 105 (3), 993–1029.
- Daruich, Diego and Julian Kozlowski**, “Macroeconomic Implications of Uniform Pricing,” *American Economic Journal: Macroeconomics*, July 2023, 15 (3), 64–108.
- Davis, Steven J., Cheryl Grim, John Haltiwanger, and Mary Streitwieser**, “Electricity Unit Value Prices and Purchase Quantities: U.S. Manufacturing Plants, 1963-2000,” *The Review of Economics and Statistics*, October 2013, 95 (4), 1150–1165.
- DellaVigna, Stefano and Matthew Gentzkow**, “Uniform Pricing in U.S. Retail Chains,” *The Quarterly Journal of Economics*, 2019, 134 (4), 2011–2084.
- Dhyne, Emmanuel, Amil Petrin, Valerie Smeets, and Frederic Warzynski**, “Theory for Extending Single-Product Production Function Estimation to Multi-Product Settings,” Working Paper 30784, National Bureau of Economic Research December 2022.
- , **Ayumu Ken Kikkawa, and Glenn Magerman**, “Imperfect Competition in Firm-to-Firm Trade,” *Journal of the European Economic Association*, 2022, 20 (5), 1933–1970.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Xu**, “Competition, Markups and the Gains from International Trade,” 2012. Working Paper New York University.

- , —, and **Daniel Yi Xu**, “How Costly Are Markups?,” *Journal of Political Economy*, 2023, 131 (7), 1619–1675.
- Fontaine, François, Julien Martin, and Isabelle Mejean**, “Price discrimination within and across EMU markets: Evidence from French exporters,” *Journal of International Economics*, 2020, 124 (C).
- Fontaine, Francois, Julien Martin, and Isabelle Mejean**, “Frictions and Adjustments in Firm-to-Firm Trade,” Working Papers 23-03, Chair in macroeconomics and forecasting, University of Quebec in Montreal’s School of Management May 2023.
- Grennan, Matthew**, “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review*, February 2013, 103 (1), 145–77.
- Hsieh, Chang-Tai and Peter J. Klenow**, “Misallocation and Manufacturing TFP in China and India,” *The Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- Huneus, Federico**, “Production Network Dynamics and the Propagation of Shocks,” Working Paper 2020.
- Ignatenko, Anna**, “Countervailing Power of Firms in International Trade,” Working Paper January 2023.
- Kaplan, Greg and Guido Menzio**, “The morphology of price dispersion,” *International Economic Review*, 2015, 56 (4), 1165–1206.
- , —, **Leena Rudanko, and Nicholas Trachter**, “Relative Price Dispersion: Evidence and Theory,” *American Economic Journal: Microeconomics*, August 2019, 11 (3), 68–124.
- Koike-Mori, Yasutaka and Antonio Martner**, “Aggregating Distortions in Networks with Multi-Product Firms,” *Working paper*, October 2025.
- Loecker, Jan De, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik**, “Prices, markups, and trade reform,” *Econometrica*, 2016, 84 (2), 445–510.
- Lorenzini, Luca and Antonio Martner**, “Aggregate Outcomes of Nonlinear Prices in Supply Chains,” Working Paper, UCLA 2025.
- Macedoni, Luca and Elena Mattana**, “Pricing in Firm-to-Firm Trade. Evidence from a Danish Multinational,” *Review of World Economics*, July 2023.
- Marshall, Guillermo**, “Search and Wholesale Price Discrimination,” *RAND Journal of Economics*, June 2020, 51 (2), 346–374.

Osotimehin, Sophie and Latchezar Popov, “Misallocation and intersectoral linkages,” *Review of Economic Dynamics*, 2023.

Pellegrino, Bruno, “Product Differentiation and Oligopoly: a Network Approach,” *Working paper*, February 2023.


Peters, Michael, “Heterogeneous Markups, Growth, and Endogenous Misallocation,” *Econometrica*, September 2020, 88 (5), 2037–2073.

Sangani, Kunal, “Markups Across the Income Distribution: Measurement and Implications,” *Working Paper* December 2025.

Appendix


A Additional tables and figures

Figure A.1: Distribution of price gaps across buyer-products



Notes: The figure shows the histogram of price gaps $\mu_{ibt} \equiv \log(p_{ibt}/p_{it})$ for $t = \text{June 2023}$. The average product price $p_{it} \equiv \sum_b s_{ib} p_{ibt}$ is computed as the quantity weighted average across buyers. Each product is weighted by its sales share. Each buyer within each product is weighted by its quantity share in the product, s_{ib} .

Figure A.2: Price dispersion and inflation across sectors




Notes: Each marker represents a sector-month. The x-axis plots the year-to-year change in producer prices for the sector-month. The y-axis plots the price dispersion for the sector-month.

Table A.1: Dispersion of price gaps: robustness for interdecile and interquartile ranges

	Unweighted		Transactions		Long codes		EAN	
	idr_i	iqr_i	idr_i	iqr_i	idr_i	iqr_i	idr_i	iqr_i
P1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P25	0.07	0.02	0.04	0.00	0.00	0.00	0.08	0.01
P50	0.16	0.08	0.14	0.06	0.09	0.02	0.18	0.09
P75	0.28	0.15	0.26	0.15	0.22	0.10	0.28	0.18
P90	0.46	0.27	0.40	0.25	0.38	0.24	0.45	0.28
P95	0.58	0.38	0.52	0.34	0.50	0.36	0.60	0.41
P99	0.93	0.67	0.86	0.64	0.93	0.69	1.00	0.65
Mean	0.21	0.12	0.18	0.10	0.15	0.08	0.21	0.12

Notes: idr_i and iqr_i respectively denote the (quantity weighted) interdecile range and interquartile range for product i . Column 'Unweighted' uses unweighted standard deviations. Column 'Transactions' uses transaction level price gaps μ_{ibr} instead of aggregating at the buyer level. Column 'Long Codes' uses the sample of products with codes that are at least eight characters in length. Column 'EAN' uses the sample of products that can be identified by an EAN code. The distributions only include products with multiple buyers. P1-P99 represent percentiles of the sales weighted distribution.

Figure A.3: Persistence in price gaps conditional on a price change



Notes: The left panel plots the coefficients β^k from equation (5) for a subset of product-buyer pairs that exhibit non-zero price changes between t and $t - k$, estimated by OLS, and for $k = 1, \dots, 12$. The regression includes product fixed effects, and observations are weighted by sales. The red lines represent 95 % confidence intervals. The right panel plots the corresponding partial R-squared of the regressions.

Table A.2: Dispersion of price gaps: low dispersion samples

	Multi-establishment buyers			Other sectors		
	Baseline	Across buyers	Within buyers across est.	Oil products	Grocery stores	Fuel retail
<i>Panel (a): Interdecile range</i>						
p5	0.00	0.00	0.00	0.00	0.00	0.00
p10	0.00	0.00	0.00	0.00	0.00	0.00
p25	0.02	0.00	0.00	0.01	0.00	0.01
p50	0.12	0.06	0.00	0.01	0.02	0.02
p75	0.24	0.22	0.00	0.01	0.11	0.05
p90	0.38	0.34	0.04	0.01	0.24	0.17
p95	0.49	0.47	0.08	0.01	0.32	0.23
Mean	0.16	0.13	0.01	0.01	0.08	0.05
<i>Panel (b): Interquartile range</i>						
p5	0.00	0.00	0.00	0.00	0.00	0.00
p10	0.00	0.00	0.00	0.00	0.00	0.00
p25	0.00	0.00	0.00	0.01	0.00	0.00
p50	0.04	0.01	0.00	0.01	0.00	0.01
p75	0.13	0.11	0.00	0.01	0.04	0.02
p90	0.23	0.23	0.01	0.01	0.13	0.11
p95	0.34	0.31	0.04	0.01	0.21	0.14
Mean	0.09	0.07	0.01	0.01	0.04	0.03
<i>Panel (c): Standard deviation</i>						
p5	0.00	0.00	0.00	0.00	0.00	0.00
p10	0.01	0.00	0.00	0.00	0.00	0.00
p25	0.03	0.00	0.00	0.01	0.00	0.01
p50	0.06	0.05	0.00	0.01	0.02	0.01
p75	0.11	0.10	0.00	0.01	0.05	0.03
p90	0.16	0.16	0.03	0.01	0.10	0.07
p95	0.21	0.19	0.05	0.03	0.13	0.09
Mean	0.08	0.06	0.01	0.01	0.04	0.03

Notes: The table reports the sales-weighted distribution of the interdecile range, interquartile range and standard deviation of price gaps across buyers of product i . In each panel, column Baseline reports the distributions of idr_i , iqr_i and σ_i respectively, and correspond to Columns 1-3 in Table 4 in the main text. The Columns 'Multi-establishment buyers' reports standard deviations for the subsample of 12,807 multi-buyer products that are sold to buyers that have multiple establishments. This subsample consists of 995 buyers. 'Across-buyers' reports the standard deviation across buyers, and 'Within buyers across est.' reports standard deviations across establishments that belong to the same buyer. Column 'Oil Products' report standard deviations for products sold by Oil producers to other firms. 'Grocery stores' and 'Fuel Retail' report standard deviations for products sold by grocery stores and fuel retailers to firms. The distributions only include products with multiple buyers. P1-P99 represent percentiles of the distribution.

Table A.3: Price dispersion by sector

σ_i													
	Wood				Pharma-			Non-metallic		Metal		Machinery	
	Food	Beverages	Textiles, Shoes	Wood prod-ucts	Paper	Printing	Chemicals	Pharma-ceutical	Plastic, Rubber	Non-metallic minerals	Metal prod-ucts	Electric equip	Machinery and equip
P1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
P10	0.01	0.00	0.00	0.00	0.01	0.01	0.01	0.00	0.00	0.02	0.01	0.01	0.02
P25	0.03	0.03	0.02	0.02	0.02	0.03	0.03	0.01	0.02	0.04	0.03	0.02	0.03
P50	0.06	0.07	0.05	0.05	0.09	0.05	0.08	0.04	0.05	0.09	0.04	0.05	0.07
P75	0.10	0.11	0.08	0.06	0.17	0.09	0.14	0.11	0.11	0.10	0.07	0.09	0.10
P90	0.14	0.16	0.13	0.09	0.21	0.16	0.23	0.21	0.18	0.13	0.12	0.14	0.16
P95	0.17	0.20	0.17	0.11	0.24	0.24	0.28	0.28	0.22	0.16	0.15	0.16	0.21
P99	0.30	0.28	0.31	0.26	0.36	0.38	0.41	0.44	0.32	0.27	0.21	0.26	0.35
Mean	0.07	0.08	0.06	0.05	0.10	0.08	0.10	0.08	0.07	0.08	0.06	0.06	0.07
Share multibuyer	0.58	0.81	0.57	0.50	0.53	0.38	0.57	0.75	0.38	0.66	0.29	0.59	0.49

Notes: The table reports the distribution of the standard deviation of prices across buyers of the same product within the different manufacturing subsectors in our baseline sample. The distributions are based on the subsample of products that were purchased by more than one buyer in June 2023. 'Share multibuyer' refers to the share of subsector's sales that is accounted for by products that are purchased by more than one buyer.

Table A.4: Persistence in price gaps

	$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$	$k = 6$
β_k	0.859*** (0.0102)	0.821*** (0.0103)	0.772*** (0.0131)	0.738*** (0.0160)	0.752*** (0.0132)	0.714*** (0.0148)
Partial R2	0.716	0.633	0.606	0.568	0.560	0.513
Observations	1,918,441	1,783,233	1,866,622	1,570,371	1,545,399	1,248,038

	$k = 7$	$k = 8$	$k = 9$	$k = 10$	$k = 11$	$k = 12$
β_k	0.692*** (0.0168)	0.686*** (0.0227)	0.680*** (0.0188)	0.665*** (0.0169)	0.618*** (0.0150)	0.618*** (0.0157)
Partial R2	0.489	0.487	0.466	0.460	0.426	0.425
Observations	1,235,935	1,192,064	1,120,737	1,068,876	1,012,200	924,888

Notes: The table reports the results depicted in Figure 2, which correspond to estimating equation (5) by OLS separately for $k = 1, \dots, 12$. All columns include product-level fixed effects.

Table A.5: Price gaps and characteristics of the buyer-seller relation

	(1)	(2)	(3)	(4)	(5)
$\mathbb{I}(\text{consistent}_{bs})$	-0.0033*** (0.0004)	0.0008 (0.0004)			
$\mathbb{I}(\text{buyer buys from seller}_{bs})$			0.0138*** (0.0030)		0.0292*** (0.0059)
log number of products _{bs}				-0.0134*** (0.0012)	-0.0092* (0.0041)
log number of transactions _{bs}				0.0068*** (0.0010)	0.0206*** (0.0042)
log sales _{s(i)b}		-0.0053*** (0.0004)			-0.0162*** (0.0029)
$0 < \text{relation}_{s(i)b} < 1 \text{ yr}$					-0.0111*** (0.0031)
$\text{relation}_{s(i)b} \geq 1 \text{ yr}$		-0.0033 (0.0022)			-0.0126*** (0.0035)
Product FE	Y	Y	Y	Y	Y
Buyer's municipality FE	Y	Y	Y	Y	Y
Buyer's sector. FE	Y	Y	Y	Y	Y
Partial R2	0.0440	0.0495	0.0414	0.0429	0.0586
Observations	3,142,733	3,142,733	3,549,623	3,549,623	3,549,623

Notes: Column 1 reports the results of a regression of log price on a dummy indicating whether the buyer made purchases from the seller in every month between June 2023-May 2024, estimated on the sample of buyer-sellers that have relations that are at least one year old. Columns 2 add controls for total buyer-seller sales and relation length. Column 3 estimates a regression of log price on a dummy indicating whether the seller made purchases from the buyer. Column 4 estimates log prices on the log of the number of number of transactions and the log of the number of products transacted between the buyer and the seller. Standard errors in parentheses are clustered at the buyer level * significant at the 10% level, **: significant at the 5% level, *** significant at the 1% level.

Table A.6: Fitted distribution of price gaps

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\sigma_i(\hat{\mu}_i)$							
Median	0.05	0.00	0.00	0.05	0.01	0.02	0.03
Mean	0.05	0.00	0.00	0.06	0.01	0.02	0.03
$idr_i(\hat{\mu}_i)$							
Median	0.10	0.01	0.01	0.10	0.02	0.03	0.06
Mean	0.11	0.01	0.01	0.12	0.02	0.04	0.08

Notes: For this table, we first compute fitted values, $\hat{\mu}_i$, using the estimated coefficients in Tables 8 and 9 in the main text. We then compute $\sigma_i(\hat{\mu}_i)$ and $idr_i(\hat{\mu}_i)$ for each product i and report their (sales weighted) means and median values.

Each column computes $\hat{\mu}_i$ from a different specification, as follows. Columns (1) and (4) use price gaps fitted using the buyer fixed effects and buyer-seller fixed effects estimated in Columns (1) and (5) of Table 8, respectively. Columns (2), (3), and (5) predict price gaps from differences in buyer sales, buyer purchases, and buyer-seller sales, using the coefficients estimated in Columns (3), (4), and (7) of Table 8, respectively. Column (6) predicts price gaps for June 2023 using observed quantities from June 2023 and the estimated quantity coefficient in Column (1) of Table 9. Column (7) predicts price gaps from a regression that incorporates buyer size, relationship length, sales, and buyer-product quantities, while allowing for sector-specific linear and quadratic terms.

Table A.7: Summary statistics on network

Number of products	237,468
Number of firms	128,748
Possible product-buyer pairs	> 30 billion
Product-buyer pairs with positive sales	> 3 million
Sales of $i \in S$ / total sales	0.31
Intermediate sales of $i \in S$ / total sales	0.119
Intermediate sales of i with multiple buyers / total sales	0.071

Notes: The table summarizes the network of products and firms used for our baseline counterfactual. $i \in S$ refers to products for which product codes are available.

B Data

B.1 Additional datasets

Here we outline additional datasets we use in the paper:

Form 22 (F22 or Income Declaration Form): This is a form submitted by firms to compute their annual tax payments, and therefore contains total sales in the year, intermediate inputs purchases, salaries, among other variables. We use this dataset for regressions where we measure buyer's size with their total sales, and for calibration of the model.

Affidavit 1887 (DJ1887): This is an annual form submitted by firms, where they declare all their employees with either a full-time or a part-time job contract. For each firm, it contains what months each worker was employed, her net salaries and whether it was a full-time or a part-time job.

Affidavit 1879 (DJ1879): This is an annual form submitted by firms, where they declare all their contracted or short-term workers. For each firm, it contains how many months each worker was employed, and how much it was retained for taxes and social security.

We use DJ1887 and DJ1879 for two purposes in the paper. First, our baseline sample uses firms (either as buyers or sellers) that have at least one paid worker in 2023. Second, for regressions that use a measure of employment at the firm level, we sum across all workers $\times$ months in both forms and divide them by 12.

B.2 Data cleaning

We take the following steps to organize and clean the transaction data in EI:

- i. We drop any transaction whose value is above 100 billion CLP (10^{10}), and price or quantity is non-positive.
- ii. We discard invoices where the seller and buyer ID are the same
- iii. We eliminate invoices that are cancelled or deemed invalid. For this, we use credit notes which give *new* or *updated* total amounts for each invoice. We drop documents where the new total amount differs by more than 1% with respect to the original total amount in the invoice
- iv. We keep firms that in 2023 had at least one paid worker, and a valid F22.
- v. Then, we define a product using the seller ID, the product code and the municipality for the seller declared in an invoice. Firms can report up to five codes, and whenever they declare the code is EAN we use it. Firms write the municipality, thus it is open to typos. When we cannot identify a selling municipality, we define the product only using the seller ID and product code.
- vi. To further refine our product definition, we only consider codes that fulfill the following: (i) are at least 4 characters long, (ii) have a number, (iii) for a given text description of the product, there is a unique product code, and (iv) for a given product code, there is a unique text description of the product.

- vii. We eliminate transactions where the product's recorded price deviates from the average price recorded in other transactions from the same month by more than 1 log point.

C Model derivations

C.1 Equilibrium characterization

The cost of the input bundle for product i is given by

$$mc_i = \left[[1 - \bar{\alpha}_i] w^{1-\sigma} + \bar{\alpha}_i q_i^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (\text{C.1})$$

Here w is the cost of the fixed factors, and q_i is the cost of the intermediate input bundle:

$$q_i = \left[\sum_j \bar{\omega}_i^j [q_i^j]^{1-\eta} \right]^{\frac{1}{1-\eta}}, \quad (\text{C.2})$$

with

$$q_i^j = \left[\sum_{i' \in j} \bar{\omega}_{i'}^j p_{i'}^{1-\rho} \right]^{\frac{1}{1-\rho}}. \quad (\text{C.3})$$

The price of the final good is:

$$P = \left[\sum_j \gamma^j [P^j]^{1-\eta} \right]^{\frac{1}{1-\eta}}, \quad (\text{C.4})$$

with

$$P^j = \left[\sum_{i \in j} \bar{\gamma}_i^j p_{ic}^{1-\rho} \right]^{\frac{1}{1-\rho}}. \quad (\text{C.5})$$

Demands are given by:

$$x_{ii'} = \bar{\omega}_{ii'} \left[\frac{p_{ii'}}{q_{i'}^j} \right]^{-\rho} m_{i'}^j; \quad c_i = \bar{\gamma}_i \left[\frac{p_{ic}}{P^{j(i)}} \right]^{-\rho} C^{j(i)}. \quad (\text{C.6})$$

with

$$\begin{aligned} m_i^j &= \bar{\omega}_i^j \left[\frac{q_i^j}{q_i} \right]^{-\eta} m_i, \\ m_i &= \bar{\alpha}_i \left[\frac{q_i}{mc_i} \right]^{-\sigma} b_i, \end{aligned} \quad (\text{C.7})$$

and

$$C^j = \bar{\gamma}^j \left[\frac{P^j}{P} \right]^{-\eta} C. \quad (\text{C.8})$$

Demand for the fixed factor is

$$l_i = [1 - \bar{\alpha}_i] \left[\frac{w}{mc_i} \right]^\sigma y_i. \quad (\text{C.9})$$

An equilibrium for this economy is given by product level prices, $\{p_{ii}\}$, $\{p_{ic}\}$, costs $\{mc_i, q_i\}$, $\{q_i^j\}$, quantities $\{c_i\}$, $\{y_i\}$ and inputs choices: $\{l_i\}$, $\{m_i\}$, $\{m_i^j\}$, and $\{x_{ii'}\}$, final prices $\{P^j\}$, P and bundles $\{C^j\}$, C and factor prices: such that: $w = 1$ and the pricing equations (9), (C.1), (C.2), (C.3), (C.4), (C.5), the demands (C.6), (C.7), (C.8), (C.9), and market clearing (10), (11) are satisfied.

C.2 Derivation of Equation (19)

This Section derives (19) in the text. Equations (16) and (17) imply

$$1 = \sum_i s_i^l \widehat{mc}_i^\sigma \left[\frac{c_i}{y_i} \hat{p}_{ic}^{-\rho} [\hat{P}^j]^{\rho-\eta} \hat{P}^\eta \hat{C}^b + \left[1 - \frac{c_i}{y_i} \right] \sum_{i'} s_{ii'} \hat{p}_{ii'}^{-\rho} [\hat{q}_{i'}^j]^{\rho-\eta} \hat{q}_{i'}^{\eta-\sigma} \widehat{mc}_{i'}^\sigma \hat{y}_{i'} \right]$$

Under the assumptions that $\alpha_i = 0$ and $c_i/y_i = 0$ for $i \in N_u$ and $\alpha_i = 1$ and $c_i/y_i = 1$ for $d \in N_d$, this simplifies to

$$1 = \sum_{i \in N_u} s_i^l \left[\sum_{i' \in N_d} s_{ii'} \hat{p}_{ii'}^{-\rho} \hat{P}^\rho \hat{C}^b \right],$$

Substituting for $\hat{P}$ and $\hat{p}_{ic}$ and rearranging yields

$$\hat{C}^b = \left[\sum_{i \in N_u} s_i^l \left[\sum_{i' \in N_d} s_{ii'} \hat{p}_{ii'}^{-\rho} \right] \right]^{-1} \left[\sum_{i \in N_d} s_i^c \hat{mc}_i^{1-\rho} \right]^{\frac{-\rho}{1-\rho}}.$$

Finally, substituting $\hat{p}_{ii'}$ for $i \in N_u$ and for $\hat{m}c_i$ for $i \in N_d$ we obtain

$$\hat{C}^b = \left[\sum_{i \in N_u} s_i^l \sum_{i' \in N_d} s_{ii'} \hat{\mu}_{ii'}^{-\rho} \right]^{-1} \left[\sum_{i \in N_d} s_i^c \sum_{i' \in N_u} \omega_{i'i} \hat{\mu}_{i'i}^{1-\rho} \right]^{\frac{\rho}{\rho-1}}.$$

Note that for $i \in N_u$, $s_i^l = \frac{wl_i}{wL} = \frac{\bar{\mu}_u}{\bar{\mu}_i} \frac{\sum_{i' \in N_d} p_{ii'} x_{ii'}}{\sum_{i' \in N_u} \sum_{i'' \in N_d} p_{ii'} x_{ii'}} = \frac{\bar{\mu}_u}{\bar{\mu}_i} s_i^x$, where $\bar{\mu}_u \equiv \frac{\sum_{i \in N_u} \sum_{i' \in N_d} p_{ii'} x_{ii'}}{wL} = \left[\sum_{i \in N_u} \bar{\mu}_i^{-1} s_i^x \right]^{-1}$. For $i \in N_d$:

$$\begin{aligned} s_i^c \omega_{i'i} \hat{\mu}_{i'i} &= \frac{p_{ic} c_i}{\sum_{i \in N_d} p_{ic} c_i} \frac{p_{i'i} x_{i'i}}{\sum_{i' \in N_u} p_{i'i} x_{i'i}} \frac{\bar{\mu}_{i'}}{\bar{\mu}_i} \\ &= \frac{p_{ic} c_i}{\sum_{i' \in N_u} p_{i'i} x_{i'i}} \frac{\sum_{i \in N_d} p_{i'i} x_{i'i}}{\sum_{i \in N_d} p_{ic} c_i} \frac{\sum_{i \in N_d} p_{i'i} x_{i'i}}{\sum_{i' \in N_u} \sum_{i \in N_d} p_{i'i} x_{i'i}} \frac{p_{i'i} x_{i'i}}{\sum_{i \in N_d} p_{i'i} x_{i'i}} \frac{\bar{\mu}_{i'}}{\bar{\mu}_i} \\ &= \frac{\bar{\mu}_i}{\bar{\mu}_d} s_{i'}^x s_{i'i}, \end{aligned}$$

where $\bar{\mu}_d \equiv \frac{\sum_{i \in N_d} p_{ic} c_i}{\sum_{i \in N_d} \sum_{i' \in N_u} p_{i'i} x_{i'i}} = \left[\sum_{i \in N_d} \bar{\mu}_i^{-1} s_i^c \right]^{-1}$. Substituting above and rearranging yields equation (19) in the text.

Finally, using the notation $\xi_i \equiv \sum_{i' \in N_d} \hat{\mu}_{ii'}^{-\rho} s_{ii'}$ to summarize the cross-buyer dispersion in markups for product i , and assuming that ξ_i is uncorrelated with $\bar{\mu}_i$, we can write the first bracket in equation (19) as

$$\sum_{i \in N_u} \frac{\bar{\mu}_u}{\bar{\mu}_i} \xi_i s_i^x = \sum_{i \in N_u} \frac{\bar{\mu}_u}{\bar{\mu}_i} s_i^x \times \sum_{i \in N_u} \xi_i s_i^x = \sum_{i \in N_u} \xi_i s_i^x. \quad (\text{C.10})$$

In addition, if the wedges are uncorrelated with the average markups set by the buyers, $\sum_{i' \in N_d} \bar{\mu}_{i'} \hat{\mu}_{ii'}^{-\rho} s_{ii'} = \sum_{i' \in N_d} \bar{\mu}_{i'} s_{ii'} \xi_i$, we can write the second bracket in (19) as

$$\sum_{i \in N_u} s_i^x \sum_{i' \in N_d} \frac{\bar{\mu}_{i'}}{\bar{\mu}_d} \hat{\mu}_{ii'}^{-\rho} s_{ii'} = \sum_{i \in N_u} \left[\xi_i \sum_{i' \in N_d} \frac{\bar{\mu}_{i'}}{\bar{\mu}_d} s_{ii'} \right] s_i^x. \quad (\text{C.11})$$

Finally, assuming that $\sum_{i \in N_u} [\xi_i \sum_{i' \in N_d} \bar{\mu}_{i'} s_{ii'}] s_i^x = [\sum_{i \in N_u} \xi_i s_i^x] \times [\sum_{i \in N_u} \sum_{i' \in N_d} \bar{\mu}_{i'} s_{ii'} s_i^x] = [\sum_{i \in N_u} \xi_i s_i^x] \bar{\mu}_d$, and substituting (C.10) and (C.11) into (19) we obtain expression (20) in the text.

D Calibration details

D.1 Apportioning firm-level inputs to products

Here we show how to apportion firm-level inputs to products under our assumptions in Section 5. First, note that if firm's uses the same input bundle across its multiple products,

then the share of input i' in the production of i is equal to the share of input i in firm's $f(i)$ total cost:

$$\frac{p_{i'i}x_{i'i}}{\sum_{i \in f(i)} p_{i'i}x_{i'i}} = \frac{\sum_{i'} p_{i'i}x_{i'i}}{\sum_{i \in f(i)} \sum_{i'} p_{i'i}x_{i'i}} = \frac{wl_i + \sum_{i'} p_{i'i}x_{i'i}}{\sum_{i \in f(i)} [wl_i + \sum_{i'} p_{i'i}x_{i'i}]}$$

If in addition the firm sets the same average markup across all its products, then this share is equal to the share of product i in firm's $f(i)$ sales.

$$\frac{wl_i + \sum_{i'} p_{i'i}x_{i'i}}{\sum_{i \in f(i)} [wl_i + \sum_{i'} p_{i'i}x_{i'i}]} = \frac{p_i^c c_i + \sum_{i'} p_{ii'}x_{ii'}}{\sum_{i \in f(i)} [p_i^c c_i + \sum_{i'} p_{ii'}x_{ii'}]}.$$

Finally, if the share of sales that go to final consumers is the same for all products of the same firm, we obtain

$$\frac{p_i^c c_i + \sum_{i'} p_{ii'}x_{ii'}}{\sum_{i \in f(i)} [p_i^c c_i + \sum_{i'} p_{ii'}x_{ii'}]} = \frac{\sum_{i'} p_{ii'}x_{ii'}}{\sum_{i \in f(i)} \sum_{i'} p_{ii'}x_{ii'}} \equiv \kappa_i.$$

Thus, $\frac{p_{i'i}x_{i'i}}{\sum_{i \in f(i)} p_{i'i}x_{i'i}} = \frac{\sum_{i'} p_{i'i}x_{i'i}}{\sum_{i'} \sum_{i \in f(i)} p_{i'i}x_{i'i}} = \kappa_i$. We can then obtain the equations in the text:

$$s_{ii'} \equiv \frac{x_{ii'}}{\sum_b x_{ib}} = \frac{x_{if(i')}}{\sum_b x_{ib}} \times \frac{x_{ii'}}{x_{if(i')}} = \frac{x_{if(i')}}{\sum_{i'} x_{ii'}} \kappa_i.$$

and

$$s_i^m = \frac{\sum_{i'} p_{i'i}x_{i'i}}{\sum_{i'} \sum_b p_{i'b}x_{i'b}} = \frac{\sum_{i'} \sum_{i \in f(i)} p_{i'i}x_{i'i}}{\sum_{i'} \sum_b p_{i'b}x_{i'b}} \times \frac{\sum_{i'} p_{i'i}x_{i'i}}{\sum_{i'} \sum_{i \in f(i)} p_{i'i}x_{i'i}} = s_{f(i)}^m \kappa_i.$$

D.2 Calibrating value added shares

As described in the text, an input in the calibration is the ratio of value added to revenues for each product i , v_i . This ratio, which can be obtained from the income declaration forms, must satisfy (for positive prices and quantities) $1 - v_i = \frac{\sum_{i'} p_{i'i}x_{i'i}}{p_{ic}c_i + \sum_{i'} p_{ii'}x_{ii'}} \leq \frac{\sum_{i'} p_{i'i}x_{i'i}}{\sum_{i'} p_{ii'}x_{ii'}} = \frac{s_i^m}{s_i^x}$. We thus compute the ratio of input purchases to revenues in the income declaration

forms and set $1 - v_i = \min \left\{ \frac{\text{inputs}_{f(i)}^{IDF} \times \theta_{f(i)}}{\text{revenue}_{f(i)}^{IDF}}, 1 - \frac{s_i^m}{s_i^x} \right\}$.

E Solution Algorithm

We now present the solution algorithm.

- i. **Step 1:** Solve for all price changes. We proceed in the following steps. First, note that defining $\widehat{mc}_i^* \equiv \widehat{mc}_i^{1-\rho}$ and $\hat{\mu}_{ii'}^* \equiv \hat{\mu}_{ii'}^{1-\rho}$, the change in marginal costs can be written as:

$$\widehat{mc}_i^* = \left[1 - \alpha_i + \alpha_i \hat{q}_i^* \frac{1-\sigma}{1-\eta} \right]^{\frac{1-\rho}{1-\sigma}}.$$

with

$$\hat{q}_i^* = \left[\sum_j \omega_i^j \hat{q}_i^{j*} \right]; \quad \hat{q}_i^{j*} = \left[\frac{1}{\omega_i^j} \sum_f p_{fi}^* \right]^{\frac{1-\eta}{1-\rho}}$$

and

$$p_{fi}^* \equiv \sum_{i' \in f} \omega_{i'i} \hat{\mu}_{i'i}^* \widehat{mc}_{i'}^*.$$

Second, note that since in our calibration, buyer-level markups, factor and input shares are firm but not product specific, i.e. $\alpha_i = \alpha_{f(i)}$, $\hat{\mu}_{ii'}^* = \hat{\mu}_{if(i')}^*$, $\omega_{i'i} = \omega_{i'f(i)}$. This implies that changes in marginal costs are the same for all products of the same firm, $\widehat{mc}_i^* = \widehat{mc}_{f(i)}^*$. Thus, we can use the following algorithm to find all changes in firm level marginal costs:

- (a) **Guess change in marginal cost of each firm:** $\widehat{mc}_f^* \equiv \widehat{mc}_f^{1-\rho}$.
- (b) **Define and calculate firm-to-firm weighted markups:** $\hat{\mu}_{ff'}^* \equiv \sum_{i \in f} \omega_{if'} \hat{\mu}_{if'}^{1-\rho}$.
- (c) **Iterate the following system to find $\widehat{mc}_f^*$**

$$\hat{p}_{ff'}^* \equiv \hat{\mu}_{ff'}^* \times \widehat{mc}_f^*,$$

and

$$\hat{q}_f^* = \sum_j \omega_f^j \hat{q}_f^{j*}; \quad \hat{q}_f^{j*} = \left[\frac{1}{\omega_f^j} \sum_{f \in j} \hat{p}_{ff'}^* \right]^{\frac{1-\eta}{1-\rho}}.$$

and

$$\widehat{mc}_f^* = \left[1 - \alpha_f + \alpha_f \hat{q}_f^{*\frac{1-\sigma}{1-\eta}} \right]^{\frac{1-\rho}{1-\sigma}}.$$

(d) **Compute all prices:**

$$\begin{aligned}\widehat{mc}_f &= \left[\widehat{mc}_f^* \right]^{\frac{1}{1-\rho}} \\ \hat{p}_{if'} &= \hat{\mu}_{if'} \widehat{mc}_{f(i)} \\ \hat{p}_{ic} &= \hat{\mu}_{f(i)c} \widehat{mc}_{f(i)} \\ \hat{q}_{f'}^j &= \left[\hat{q}_{f'}^{j*} \right]^{\frac{1}{1-\eta}} \\ \hat{q}_{f'} &= \left[\hat{q}_{f'}^* \right]^{\frac{1}{1-\eta}} \\ \hat{p}^j &= \left[\frac{1}{s^{jc}} \sum_{i \in j} s_i^c \hat{p}_{ic}^{1-\rho} \right]^{\frac{1}{1-\rho}}\end{aligned}$$

and

$$\hat{p} = \left[\sum_j s^{jc} \left[\hat{p}^j \right]^{1-\eta} \right]^{\frac{1}{1-\eta}}.$$

ii. **Step 2:** Solve for changes in the total quantity of each product relative to the change in final consumption, $\frac{\hat{y}_i}{\hat{C}}$. We proceed in the following steps.

(a) **Compute change in relative final demand of each product:**

$$\frac{\hat{c}_i}{\hat{C}} = \left[\frac{\hat{p}_{ic}}{\hat{p}^{k(i)}} \right]^{-\rho} \left[\frac{\hat{p}^{k(i)}}{\hat{p}} \right]^{-\eta}.$$

(b) **For products that are only sold to final consumers, set**

$$\frac{\hat{y}_i}{\hat{C}} = \frac{\hat{c}_i}{\hat{C}}.$$

(c) **Then guess $\frac{\hat{y}_i}{\hat{C}}$ for products $[1 : NP_{XT}]$, compute**

$$\frac{\hat{b}_f}{\hat{C}} = \sum_{i \in f} \kappa_i \frac{\hat{y}_i}{\hat{C}}.$$

and update the guess using

$$\frac{\hat{y}_i}{\hat{C}} = \widehat{FDEM}_i + \sum_f INTDEM_{if} \frac{\hat{b}_f}{\hat{C}}.$$

with $\widehat{FDEM}_i \equiv \frac{c_i}{y_i} \frac{\hat{c}_i}{\hat{C}}$, and $INTDEM_{if} = \left[1 - \frac{c_i}{y_i}\right] \frac{x_{if}}{\sum_f x_{if}} \left[\frac{\hat{p}_{if}}{\hat{q}_f^{k(i)}}\right]^{-\rho} \left[\frac{\hat{q}_f^{k(i)}}{\hat{q}_f}\right]^{-\eta} \left[\frac{\hat{q}_f}{\widehat{mc}_f}\right]^{-\sigma}$

(d) **From the labor market clearing, solve:**

$$\hat{C} = \left[\sum_i s_i^l \widehat{mc}_i^\sigma \frac{\hat{y}_i}{\hat{C}} \right]^{-1}.$$