Multiproduct Transactions in International Markets

Eric W. Bond [*], Ji hye Heo[†], Ebehi Iyoha [‡], Joel Rodrigue[§]

November 6, 2026

Abstract

This paper establishes that roughly 30 percent of the value of all trade occurs within multiproduct *transactions* (MPTs). We document four empirical patterns: (i) MPTs involve few products and have small transaction sizes; (ii) among the importers that conduct MPTs, the share of transactions that are MPTs decreases with firm size; (iii) conditional on firm size, larger purchases of a product are less likely to be in an MPT; and (iv) importers engaging in MPTs source from relatively few suppliers. We propose a model of international sourcing rationalizing these features of the data. The model predicts that trade cost reductions lead to a reduction in MPTs, sourcing from a wider set of suppliers, and additional input cost reductions, reinforcing gains from trade.

Keywords: Importing, Multi-Product Firms, Multi-Product Transactions.
JEL Codes: F12, F14, L11, L14 .

[*]Department of Economics, Vanderbilt University, VU Station B #351819, 2301 Vanderbilt Place, Nashville, TN 37235-1819; Tel.: +1 615 322 2871; fax: +1 615 343 8495. Email: eric.w.bond@vanderbilt.edu.

[†]Department of Economics, Oberlin College and Conservatory, Rice Hall, 38 E. College St., Oberlin, OH 44074, Email: jheo1@oberlin.edu.

[‡]Harvard Business School, Cambridge, MA 02138. E-mail: eiyoha@hbs.edu.

[§]Department of Economics, Vanderbilt University, VU Station B #351819, 2301 Vanderbilt Place, Nashville, TN 37235-1819; Tel.: +1 615 322 2871; fax: +1 615 343 8495. Email: joel.b.rodrigue@vanderbilt.edu.

1 Introduction

A firm sourcing several inputs internationally must decide not only which suppliers to use, but also which products to purchase from each. Buying several products from one supplier can economize on the fixed costs of sourcing, but may require paying more for inputs available at lower prices elsewhere. This paper studies this trade-off through *multiproduct transactions* (MPTs), in which an importer purchases two or more products from the same exporter in a single transaction. We document the prevalence and composition of MPTs using Mexican import records and develop a model of costly sourcing and strategic pricing to examine their implications for input prices and the gains from trade.

Firm-level product scope provides an incomplete picture of how these purchases are organized. Research on multiproduct firms examines which products firms supply to each market (Bernard, Redding, and Schott, 2011), while research on global sourcing and buyer–seller relationships studies firms' access to foreign suppliers (Antràs, Fort, and Tintelnot, 2017; Bernard, Moxnes, and Ulltveit-Moe, 2018). Neither a firm's aggregate product portfolio nor its set of suppliers, however, identifies which products it purchases together in an individual transaction. We focus on this transaction-level margin, asking when importers consolidate purchases with a supplier and how their alternatives for sourcing individual products constrain the prices they pay.

Drawing on 45.7 million import transactions from 2011 to 2020, we start by constructing a novel firm-to-firm transaction-level dataset based on sourcing among Mexican firms. The records identify the products included in each transaction and their quantities and values. MPTs account for approximately 31 percent of observed import value. Although multiproduct importers account for most import value in our data, most individual transactions contain only one product. Firm-level and transaction-level product scope therefore describe distinct aspects of international sourcing.

We proceed to establish four new empirical facts about MPTs on international markets. First, although multiproduct firms trade a wide range of products, the typical MPT is small in scope: half of MPTs include only two or three products. Moreover, a product is purchases in a smaller quantity when it is bundled only with products of other HS4 headings than when the same exporter ships it alone in the same year.

Second, large firms are more likely to have at least one MPT relative to small firms, since they are more likely to buy and sell multiple products. Conditional on conducting a MPT, however, the share of an importer's transactions that are multiproduct declines with importer size, measured by annual import value. Smaller importers thus rely more heavily on MPTs even though larger importers are more likely to engage in them at least once. For example, among importers in the bottom quartile of the firm size distribution with at least one MPT, nearly 73 percent of transactions include two or more products, on average. Among comparable firms in the top quartile, the share is 41 percent.

Third, conditional on importer size, larger purchases of a product are less likely to be part of an MPT. This pattern holds in every quartile of importer size, including the largest importers.

Fourth, the number of suppliers per product rises with importer size, from 1.04 in the bottom size decile to 2.20 in the top decile. Importers in every size decile have fewer suppliers than products on average, and each supplier provides more products to larger importers. The pattern complements the

transaction-level evidence that smaller importers rely more heavily on joint purchases.

Together, these four facts highlight systematic differences in sourcing strategies across firms and transactions: small importers rely heavily on bundling, while large importers tend to separate products across suppliers. To rationalize these patterns, we develop a model of global sourcing with costly supplier search and heterogeneous supplier advantages. In our framework, importing firms optimally search for suppliers of multiple inputs on global markets. Suppliers can produce each input but have a comparative advantage in their core product. Because searching an additional supplier entails a fixed cost, firms face a trade-off between sourcing from the cheapest supplier of each input versus bundling products from fewer suppliers.

We emphasize the *transactional* dimension of international trade rather than *sourcing* for three reasons. First, the empirical patterns we document hold within individual source countries; in this sense, the features we highlight are conceptually distinct from cross-country sourcing. Second, this emphasis highlights buyer–seller interactions in global markets, which are our margin of interest. Third, it reflects the shipment data we introduce to the broader sourcing literature and the novel patterns these data reveal.

The model predicts that less productive firms will source from a narrower set of suppliers, making it optimal to purchase multiple goods from the same supplier. In contrast, larger firms, with higher quantities demanded, have stronger incentives to source each input from its lowest-cost supplier, leading to broader supplier networks and a greater prevalence of single-product transactions.

Beyond these sourcing decisions, the model highlights an important implication of multiproduct transactions. When an importer bundles purchases with a single supplier, the supplier acquires market power over peripheral products because the importer must incur an additional fixed cost to source those products elsewhere. This creates scope for strategic pricing, whereby suppliers charge higher prices for peripheral goods while still retaining the buyer. The model therefore predicts that (i) a firm with smaller demand for a peripheral product is more likely to buy it in an MPT, and an MPT is skewed toward its supplier's core product, (ii) peripheral products included in MPTs command higher prices than from a supplier for which they are the core product, although the importer spends less on them, and (iii) a rise in a supplier's trade costs lowers its share of MPTs relative to other suppliers.

Section 5 states these predictions as testable implications and describes how we test them, by comparing MPTs with separate purchases of the same products by comparable importers and by the same importer at other dates. These tests are in progress.

The strategic pricing mechanism also has important implications for the gains from trade liberalization. We develop a welfare framework that shows trade liberalization affects firms not only by changing the prices they pay directly, but also by altering the incentives to bundle purchases within multiproduct transactions or to source products separately. As trade costs fall, firms increasingly substitute away from MPTs toward single-product transactions, strengthening competitive pressure among suppliers and reducing the prices charged for peripheral products. In this sense, the model implies procompetitive gains that arise from the fact that importers source multiple inputs from multi-product exporters. The margin between multiproduct and single-product transactions therefore becomes a central channel through which trade liberalization generates additional welfare gains.

Taken together, these findings highlight that trade costs influence international trade through more than the conventional geography of sourcing. Trade costs shape not only where firms source their inputs, but also how products are allocated across suppliers and the degree of competition those suppliers face. Holding both the set of imported products and total import quantities fixed, the ability to source individual products from their lowest-cost suppliers represents an additional source of gains from trade, arising through both more efficient supplier allocation and the pro-competitive effects generated by the transition from bundled to separate sourcing.

1.1 Literature

The dominance of multiproduct firms in international trade is well established empirically (Bernard, Jensen, Redding, and Schott, 2012; Manova and Yu, 2017; Arkolakis, Ganapati, and Muendler, 2021). We build on this literature by characterizing how transaction scope yields microeconomic foundations for multiproduct trade. Whereas MPTs typically cover only a subset of the products sourced by individual importers and vary systematically across firm and transaction size, standard models of multiproduct firms largely abstract from transaction-level sourcing trade-offs.

A large share of the multiproduct trade literature examines optimal firm- and market-specific product scope and how exporters adjust to trade liberalization (Bernard, Redding, and Schott, 2010; Eckel and Neary, 2010; Bernard, Redding, and Schott, 2011; Qiu and Zhou, 2013; Timoshenko, 2015).[1] In contrast, our focus is not on the determinants of firm-level product scope. Instead, we characterize how importer sourcing across heterogeneous multiproduct exporters influences the allocation of goods across firms and markets, following the spirit of Amador and Opromolla (2013).

A complementary empirical literature evaluates the impact of trade cost reductions on product scope (Baldwin and Gu, 2009; Iacovone and Javorcik, 2010; Bernard, Redding, and Schott, 2011; Berthou and Fontagné, 2013; Lopresti, 2016). We likewise study the role of trade costs, but emphasize their impact on the geography of MPTs and develop a model that characterizes the strategic response of international transactions to changes in trade costs.

Perhaps the strand of research closest in spirit to our paper is the work on demand trade-offs within multiproduct firms. Feenstra and Ma (2008), Baldwin and Gu (2009), Eckel and Neary (2010), and Dhingra (2013) each highlight how the introduction of additional varieties may cannibalize sales of existing products. More recently, Arkolakis, Ganapati, and Muendler (2021) argue that firm-level trade patterns are consistent with substantial economies of scope in market access costs. We contribute to this discussion by examining the extent to which such implied economies of scope can be rationalized by a model of frictional search and sourcing in international markets, and by showing how multiproduct exporters can exploit market power in such settings.

In this sense, our work also directly relates to the rapidly growing literature characterizing firm-to-firm networks in international markets (Bastos and Silva, 2012; Chaney, 2014; Bernard, Moxnes, and Ulltveit-Moe, 2018; Eaton, Jinkins, Tybout, and Xu, 2022; Demir, Fieler, Xu, and Yang, 2024). We

[1] Thomas (2011) likewise develops a model of product scope across multiproduct, multinational laundry detergent manufacturers in Western Europe.

focus in particular on how importers source across potential trade partners in disparate locations, as in (Antràs, Fort, and Tintelnot, 2017; Bernard, Jensen, Redding, and Schott, 2018), and how firm-to-firm linkages respond to trade cost shocks, as in (Huneeus, 2018; Demir, Fieler, Xu, and Yang, 2024; Arkolakis and Miyauchi, 2025). This literature broadly finds that firms highly integrated in world markets participate more intensively along both intensive and extensive margins, thereby magnifying their share of aggregate trade. In contrast, we show that MPTs are not dominated by large importers; rather, it is small importers that disproportionately source products in bundles, while large importers are more likely to source products individually.

Our analysis also contributes to the literature on the gains from trade (Arkolakis, Costinot, and Rodríguez-Clare, 2012). A large literature beginning with Krugman (1979) and Markusen (1981) emphasizes that trade liberalization can generate welfare gains by increasing competitive pressure and reducing firms' market power. More recently, Atkeson and Burstein (2008) and Edmond, Midrigan, and Xu (2015) demonstrate how endogenous markups shape the magnitude and distribution of the gains from trade. We identify a complementary channel through which trade costs affect welfare. By changing the incentives to bundle products within multiproduct transactions or to source them separately across suppliers, trade liberalization alters the degree of competition faced by exporters, the prices charged for peripheral products, and the organization of firm-to-firm sourcing relationships. Consequently, the transition from multiproduct to single-product transactions becomes an additional mechanism through which lower trade costs generate welfare gains.

1.2 Roadmap

The remainder of the paper is organized as follows. Section 2 describes the data. Section 3 documents the stylized facts, while section 4 presents a model that rationalizes these empirical patterns, highlighting their implications for strategic pricing and welfare. Section 5 provides empirical tests of the model's predictions. Section 6 concludes.

2 Data

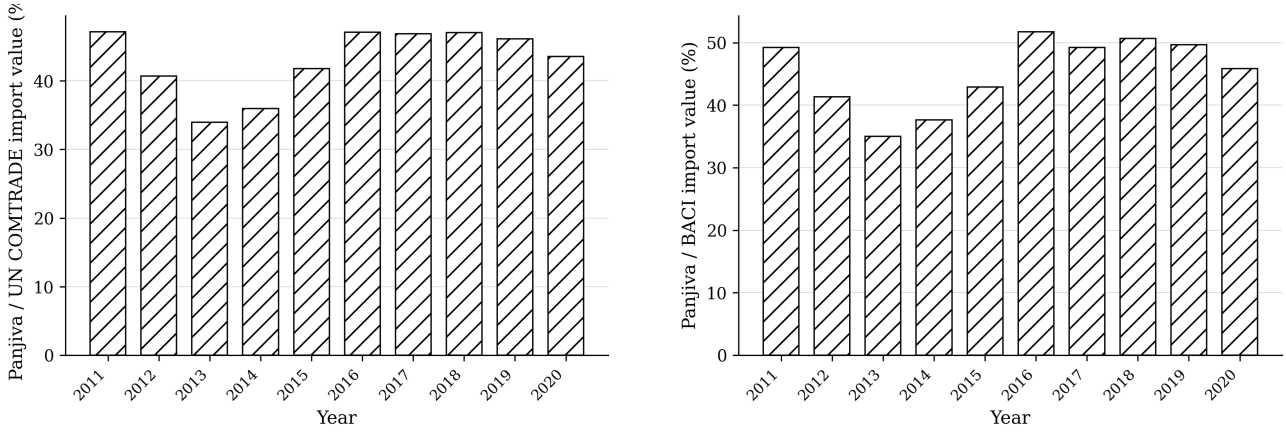
The dataset used in this paper is constructed from import transaction records curated by *Panjiva*. Each record provides transaction-level information, including the names and locations of each buyer and seller, the HS codes included in the shipment, and the quantity and value of each product shipped from seller to buyer.[2] We use all Mexican import transactions recorded by *Panjiva* between 2011 and 2020. To check the consistency of our results, we also make use of similar data curated by *Import*

[2]We measure quantity by the quantity declared for tariff purposes, in the tariff unit of the product line: kilograms for 67 percent of the liens with a usable quantity and pieces for 31 percent. Because units differ across products, we compare quantities only within a product and unit. The declared quantity is complete for 99.3 percent of product lines. We treat a quantity as missing when its value per unit lies more than 1,000 times above or below the median for its product and unit; this sets aside 0.36 percent of the lines with a complete quantity, which account for 1.7 percent of their value, and leaves a usable quantity for 99.0 percent of all product lines. With thresholds of 100 or 10,000 instaed of 1,000, the MPT share of quantity in Table 1 is 30.1 or 31.3 percent rather than 31.0 percent. No transaction is dropped for a missing quantity.

Genius.[3]

We evaluate data coverage by aggregating the transaction records by year and comparing their total value with the value of Mexican imports reported in UN COMTRADE and BACI. As shown in Figure 1, the transaction data cover a little under half of Mexican imports on average: between 34 and 47 percent of annual import value in UN COMTRADE and between 35 and 52 percent in BACI, with the lowest coverage in 2013.

Figure 1: Dataset Coverage Compared to External Sources



Notes: Annual import value in our data, after excluding transactions with missing identifiers or from exporters located in Mexico and before excluding importer–exporter pairs that transact only once, as a percentage of Mexican imports in UN COMTRADE (left) and BACI (right), 2011–2020.

2.1 Classifying Multiproduct Transactions

For our purposes, we define an *accounting MPT* as any transaction between an individual buyer (importer) and seller (exporter) that includes two or more products (HS6 codes).[4] The advantage of this definition is that it is easily identifiable in the raw data and allows us to clearly describe the nature of MPTs in international markets.

Two concerns arise with this benchmark definition. The first is that it may be too broad: a buyer and seller might contract separately on two products and then combine them ex post into a single bill of lading. In our data we cannot distinguish such cases from those in which two firms explicitly bargain and contract over multiple goods at the same time.

The second concern is that the definition may be too narrow: firms could contract simultaneously on multiple products but require that they be shipped and billed separately. This possibility is less

[3] *Panjiva* and *Import Genius* provide similar information, but with some important differences. *Panjiva* data are primarily curated from government sources and may occasionally incorporate third-party adjustments, whereas *Import Genius* provides raw data obtained via FOIA requests. Our reliance on *Panjiva* reflects the fact that the *Import Genius* dataset spans a shorter period (2015–2019) and excludes non-maritime trade. As a robustness check, we also replicate key findings using the *Import Genius* data and confirm that our conclusions are unaffected. In addition, we cross-check results with data from other countries such as Chile, Colombia, and Peru. For parsimony, however, we focus on Mexico in the main text.

[4] For consistency across all exercises, we aggregate the transaction data to the HS6 level and treat each HS6 code as a distinct product.

problematic for three reasons. First, international shipments are widely reported to be infrequent and lumpy. Indeed, in our data 45 percent of buyer–seller pairs transact only once between 2011 and 2020, and among those that transact more than once, the median number of transactions over the same period is five.[5] Second, if some true MPTs are misclassified as single-product transactions (SPTs), the observed differences between the two groups will be biased toward zero, rendering our estimates conservative. Third, our empirical exercises focus on systematic differences across transactions and firms, rather than on individual classification errors.

2.2 Sample Selection

We impose two main sample restrictions. First, we exclude transactions with missing identifiers (transaction, importer or exporter IDs), treating placeholder entries as missing.[6] Second, we drop importer exporter pairs that transact only once between 2011 and 2020. This rule also removes every firm that appears only once in the sample, which avoids artificially skewing the firm-size distribution. The first restriction removes 2.9 percent of transactions and the second removes 2.0 percent of those that remain, leaving 45.7 million transactions. Firm size, the annual value of a firm's transactions, is measured using all transactions in which the firm is identified, before these restrictions are applied. When conditioning on multiproduct sellers or buyers, we restrict the sample to firms that export or import at least two HS6 products in a given year.

2.3 Descriptive Statistics

Table 1: Multiproduct Transactions Among Mexican Imports

Total No. of Transactions	MPT Share (%) of		
	Transactions	Value	Quantity (within product)
45,678,791	31.01	31.34	31.01

Notes: Mexican import transactions, 2011–2020, after the sample restrictions of Section 2. An MPT is a transaction that includes two or more HS6 products. The first two shares are the MPT shares of the number of transactions and of import value. Quantities are not recorded in a common unit across products, so the quantity share is computed within product: the MPT share of each product's quantity in each of its tariff units, Q_p^{MPT}/Q_p, averaged across product-unit cells with import-value weights, over the product lines with a usable quantity, which hold 97.8 percent of import value. Applied to value, the same weights give an MPT share of 31.52 percent.

Table 1 documents the prevalence of MPTs in our data. Of the 45.7 million import transactions in our sample, 31 percent include two or more products. These transactions account for 31 percent of total import value.

Table 2 shows that the typical MPT includes a small number of products: the median number of products per MPT is only four, and just under half of all MPTs include two or three products. The mean is 7.3, reflecting a long right tail. The remaining rows of Table 2 split MPTs by type of seller.

[5] These statistics are computed on the 2.03 million pairs observed before we exclude pairs that transact only once (see the next subsection). The pairs that transact more than once are the pairs that remain in our sample.

[6] The placeholder identifiers are `SN`, `NA`, and `NOPROPORCIONO`.

We classify an exporter as *broad scope* in a given year if the number of HS6 products it sells exceeds the 99th percentile across exporters in that year, and as *intermediary-like* if, in addition, its largest HS2 chapter accounts for less than half of its sales. Intermediary-like sellers account for 18 percent of MPTs and 13 percent of MPT value; excluding them lowers the median to three products and the mean to 6.1.

Table 2: MPT Product Scope Among Mexican Imports

		Share of All (%)				Percentiles				Share (%) with	
	Number of MPTs	MPTs	MPT Value	Mean	Std. Dev.	25	50	75	Max	2–3 Products	One HS2 Chapter
All MPTs	14,166,326	100.00	100.00	7.29	12.28	2	4	8	786	49.98	21.62
By type of seller											
Other sellers	9,194,674	64.91	65.58	4.97	5.34	2	3	6	133	57.87	25.91
Broad scope, concentrated	2,482,655	17.53	21.53	10.39	14.22	2	5	12	398	38.60	18.00
Broad scope, intermediary-like	2,488,997	17.57	12.90	12.76	22.19	3	6	15	786	32.17	9.42
Excl. intermediary-like sellers	11,677,329	82.43	87.10	6.13	8.39	2	3	6	398	53.77	24.22

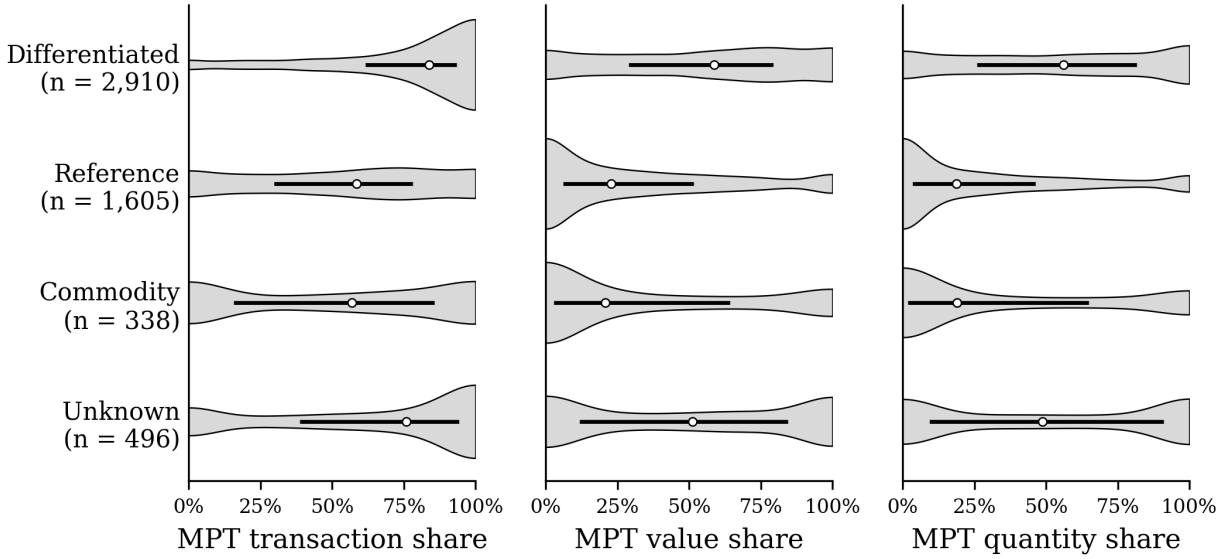
Notes: One observation per MPT, 2011–2020. Mean, standard deviation, percentiles, and maximum refer to scope, the number of distinct HS6 products in the transaction. Shares of all MPTs are by number and by import value. An exporter has broad scope in a given year if its number of HS6 products exceeds the 99th percentile across exporters in that year. A broad-scope exporter is intermediary-like if its largest HS2 chapter accounts for less than half of its sales, and concentrated otherwise. The last row excludes MPTs sold by intermediary-like exporters. The last two columns report the percentage of MPTs with two or three products and the percentage whose products all belong to one HS2 chapter.

MPTs are concentrated among differentiated products. Differentiated products account for 58 percent of total import value but 80 percent of the value imported in MPTs. Figure 2 plots the distribution across HS6 products of the MPT share of each product's transactions, value traded, and quantity shipped, by category of the conservative Rauch classification (Rauch, 1999).[7] The figure weights HS6 products equally, and an MPT enters the count of every product it contains, so these shares are not directly comparable to the aggregate shares in Table 1: MPTs make up 31 percent of transactions but 77 percent of transaction-product observations. MPTs account for a higher share of transactions, value, and quantity traded among differentiated products than among commodities or reference goods. For the median differentiated product, close to 60 percent of import value is traded in MPTs, compared with just over 20 percent for the median reference good or commodity. Moreover, in every category the distribution of the transaction share lies to the right of the corresponding distributions for value and quantity. This suggests that, within a product, the average purchase in an MPT is smaller than the average purchase in an SPT, a feature we investigate in greater detail below. Table 3 illustrates the variation across products with examples from 2017. Almost none of the import value of liquefied natural gas or of floating docks and special-function vessels is traded in MPTs, compared with 47 to 65 percent for passenger cars and light trucks and 94 percent for hot-rolled stainless steel coil.

Figure 3 reports the multiproduct share of total transactions, value traded, and kilograms shipped by the country in which the exporter is located. Each panel shows the 20 largest countries by the number

[7]Similar results are found using the liberal classification, under which differentiated products account for 55 percent of total import value and 76 percent of MPT value.

Figure 2: MPTs and Rauch Product Classification



Notes: Each observation is an HS6 product, pooled over 2011–2020; n is the number of products in each category of the conservative Rauch (1999) classification, and Unknown collects products the classification does not cover. For each product, the MPT share is the share of its transactions, import value, or quantity traded in MPTs. Shaded areas are kernel density estimates; the dot marks the median and the bar the interquartile range.

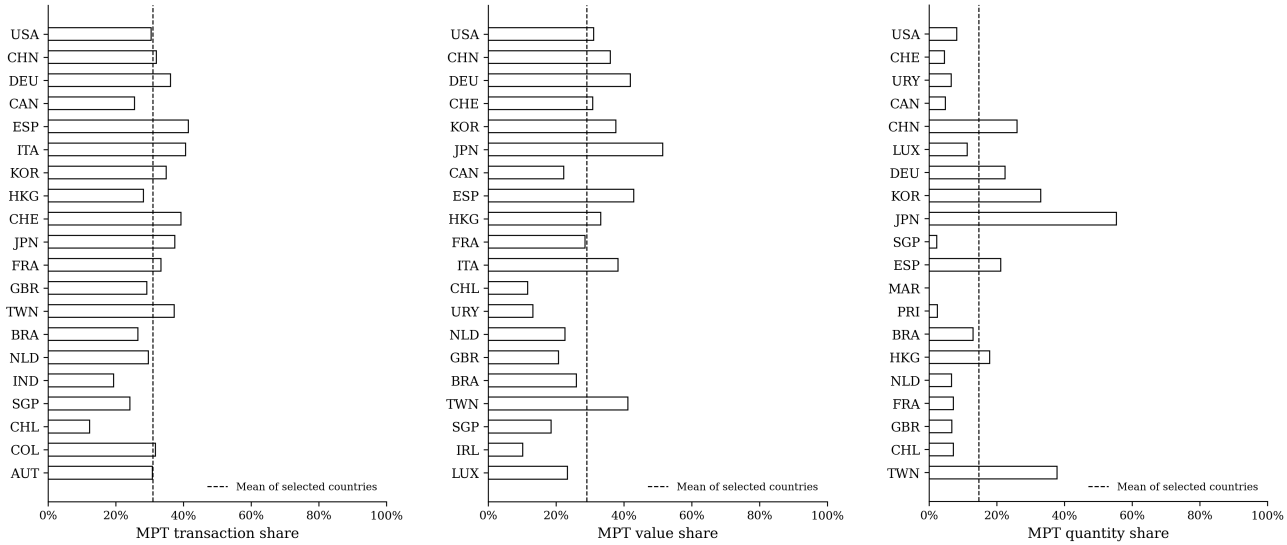
Table 3: Examples of Mexican Imports by MPT Frequency

Included in MPT?	HS Code	Description	Value (Million USD)	MPT Share of Value (%)
Rarely	271111	Natural gas, liquefied	912.91	0.00
	890590	Floating docks, special function vessels	1,118.44	0.00
Sometimes	870322	Automobiles, spark ignition, 1000–1500cc	1,838.87	64.72
	870323	Automobiles, spark ignition, 1500–3000cc*	4,114.47	47.05
	870431	Spark ignition engine trucks, <5 tonnes	1,148.25	58.99
Mostly	721913	Hot rolled stainless steel coil, w>600 t 3–4.75mm	313.17	94.09

Notes: Import value and the MPT share of import value in 2017 for six HS6 products. (*) marks a product among the five largest Mexican imports by value in 2017 in our sample.

of transactions, the value traded, or the kilograms shipped, as appropriate, sorted in descending order. Two features stand out. First, MPTs are ubiquitous: they account for between 12 and 41 percent of transactions in each of the 20 countries with the most transactions, although their shares of value and quantity vary more widely (from 10 to 51 percent of value and from close to zero to 59 percent of kilograms). Second, the MPT share of transactions does not appear to correlate strongly with source-country characteristics at the aggregate level. This suggests the importance of product-, firm-, or relationship-specific characteristics for understanding the nature of MPTs in international markets. At the same time, MPTs are distinct from multiproduct production (Bernard, Redding, and Schott, 2010, 2011; Bernard, Jensen, Redding, and Schott, 2012; Mayer, Melitz, and Ottaviano, 2014; Lopresti, 2016; Arkolakis, Ganapati, and Muendler, 2021). While 98 percent of import value in our data is purchased by multiproduct importers, only 31 percent of total import value arises from transactions that include two or more products.

Figure 3: MPTs and Country Sources

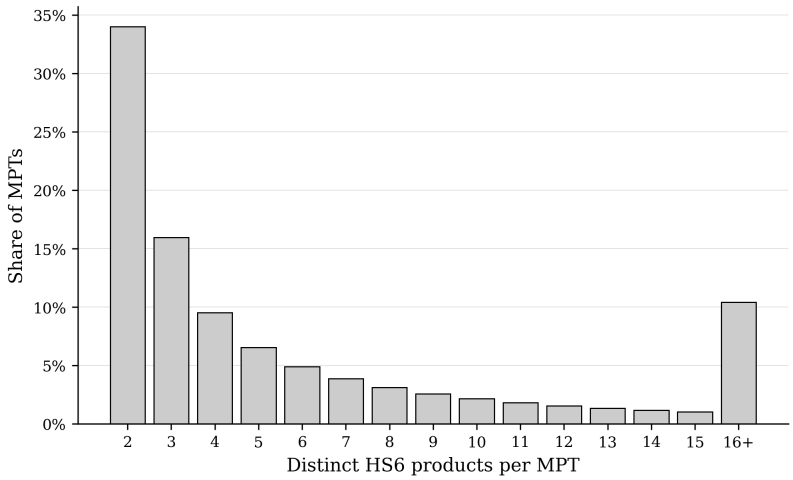


Notes: MPT share of transactions, import value, and kilograms, pooled over 2011–2020, by the country in which the exporter is located. Each panel shows the 20 largest countries by the corresponding total, in descending order, so the countries differ across panels. The dashed line is the unweighted mean of the 20 shares shown. The kilogram panel covers the product lines whose tariff unit is the kilogram and whose quantity is usable.

3 Stylized Facts

This section describes four robust empirical features of MPTs, which form the basis for our theory of multiproduct sourcing in international markets.

Figure 4: Distribution of MPT Product Scope



Notes: Each observation is an MPT in the sample of Section 2, 2011–2020 (14,166,326 MPTs). Product scope is the number of distinct HS6 products in the transaction. Bars show the share of MPTs with each scope; the last bar pools MPTs with 16 or more products.

Fact 1: The typical MPT is small in scope, and a product is purchased in a smaller quantity when it is bundled only with dissimilar products.

Figure 4 plots the distribution of MPT product scope, while Table 2 reports summary statistics. Transaction-level product scope follows a highly skewed distribution: the minimum and modal scope is two products, the median is four, and the mean is 7.3. About one third of MPTs include exactly two products, one half include two or three, and two thirds include five or fewer. The mean exceeds the median because of a long right tail: about one in ten MPTs include 16 or more products.

While it may not be surprising that the number of products per transaction is typically small, it is less obvious whether a product is purchased in a smaller quantity when it is part of an MPT than when it is purchased alone. To compare the two cases, we regress the log quantity of each product line, an HS6 product within transaction, on an indicator for lines that belong to an MPT. A second regression replaces this indicator with two: one for MPT lines whose transaction contains another product of the line's own HS4 heading, and one for MPT lines whose transaction contains none. Lines in SPTs are the omitted group. Table 4 reports the results. Column (1) includes exporter-product-year fixed effects, so it compares shipments of the same product by the same exporter in the same year. Column (2) adds importer-year fixed effects. Column (3) includes importer–exporter–product–unit–year fixed effects, so it compares purchases of the same product by the same importer from the same exporter in the same year.

On average, the quantity of a product is 10 to 11 percent smaller when it is shipped in an MPT than when the same exporter ships it alone in the same year, with or without importer-year fixed effects (Panel A). Within importer-exporter pairs the difference is 6 percent and is statistically significant only at the 10 percent level.[8] Panel B splits MPT lines by the other products in their transaction.

[8]In column (3), 1.92 million importer–exporter–product–unit–year cells contain both MPT and SPT lines. These cells

Table 4: Within-Exporter Quantity Differential by MPT Type (Mexican Imports)

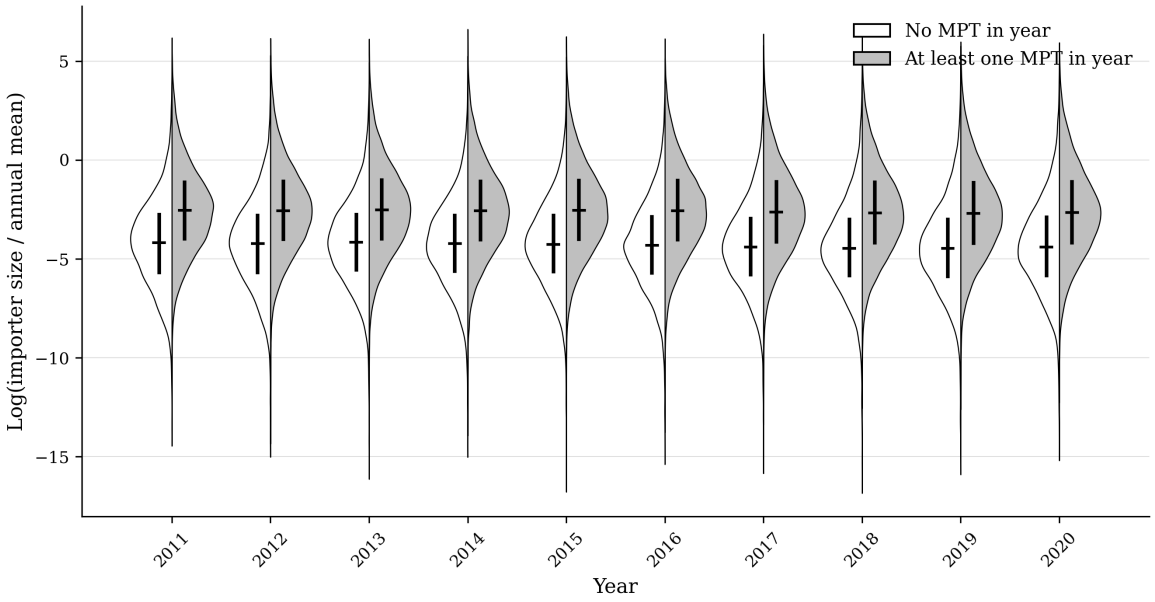
	Share of MPT lines (%)	Log quantity (1)	(2)	(3)
Panel A. All product lines in MPTs				
MPT line	100.0	-0.113^{***}	-0.107^{***}	-0.061^{*}
		(0.036)	(0.033)	(0.033)
Panel B. By whether the transaction includes another product of the line's own HS4 heading				
MPT line, no other product of its heading	63.4	-0.188^{***}	-0.178^{***}	-0.132^{***}
		(0.034)	(0.032)	(0.031)
MPT line, another product of its heading	36.6	0.077^{*}	0.079^{**}	0.130^{***}
		(0.044)	(0.040)	(0.040)
Exporter–product–year fixed effects		Yes	Yes	Nested
Importer–year fixed effects		No	Yes	Nested
Importer–exporter–product–year fixed effects		No	No	Yes
R^2, Panel A		0.863	0.869	0.884
R^2, Panel B		0.864	0.870	0.884
Observations		133,417,256	133,417,256	133,417,256

Notes: The unit of observation is a product line: an HS6 product within a transaction, with records of the same HS6 code in a transaction combined. An MPT is a transaction with two or more HS6 products; an SPT has one. The sample is the product lines with a usable quantity (99.0 percent of all product lines) of Mexican import transactions, 2011–2020, after the sample restrictions of Section 2; all products are included. The dependent variable is the log of the quantity of the line declared for tariff purposes, in its tariff unit and quantities are compared only within a unit. Each column of each panel is one regression. Lines in SPTs (23.4 percent of the lines) are the omitted group, so each coefficient is the difference in log quantity between MPT lines and SPT lines, conditional on the fixed effects of its column; a coefficient b is a difference of $100(e^b - 1)$ percent. Panel B splits MPT lines by whether their transaction contains another HS6 product of the line's own HS4 heading. Column (1) includes exporter–product–unit–year fixed effects and column (2) adds importer–year fixed effects. Column (3) includes importer–exporter–product–unit–year fixed effects, which nest the fixed effects of columns (1) and (2); in this column, 1,919,483 importer–exporter–product–year cells contain both MPT and SPT lines (41,045,962 lines). Lines that are alone in their fixed-effect cell are counted in the observations and do not affect the estimates. Standard errors clustered two-way by HS6 product (5,344 clusters) and by exporter (485,178 clusters) are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Products bundled only with products of other HS4 headings, 63 percent of all MPT lines, are 12 to 17 percent smaller than when they are shipped alone, and the difference is significant at the 1 percent level in every column. For products bundled with another product of their own heading, the estimated difference is positive and statistically significant in every column. These lines are 8 percent larger in columns (1) and (2), at the 10 and 5 percent levels, and 14 percent larger within importer–exporter pairs, at the 1 percent level[9]

Fact 2: Among importers that conduct MPTs, MPT intensity declines with firm size.

Figure 5: Firm Size Distribution Across MPT Status



Notes: Each observation is an importer in a given year in the sample of Section 2, 2011–2020 (471,504 importer-years). Importer size is the log of the importer's annual import value, computed from all transactions in which the importer is identified (Section 2), relative to the mean across importers in the same year. For each year, the left (white) density is for importers with no MPT in that year and the right (gray) density is for importers with at least one MPT. Densities are kernel estimates scaled to the same maximum width. The horizontal mark is the median and the vertical bar the interquartile range.

To characterize MPTs across the firm-size distribution, we measure importer size by the total value of a firm's imports in a given year, since the data do not report total firm sales.[10] Figure 5 compares the size distribution of importers that conduct at least one MPT in a given year with that of importers that conduct none in that year. Two features stand out. First, MPTs are common across nearly all firm sizes. Second, larger firms are more likely to have had at least one MPT in a given year. The share

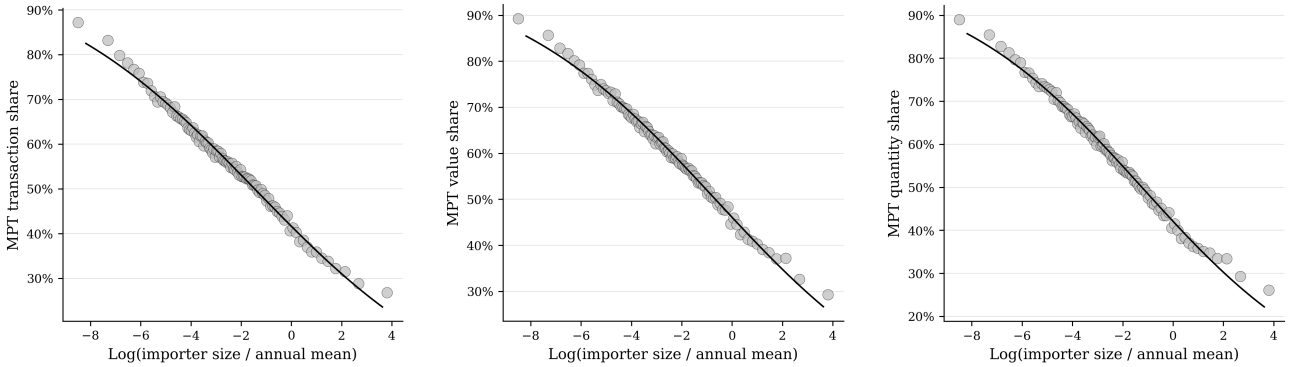
hold 41.0 million product lines, 31 percent of the lines in the regressions.

[9]Table A1 in Appendix A splits the two groups further, by the closest other product in the transaction and by whether the transaction spans several HS2 chapters. The estimate is positive and statistically significant only for lines with another product of their own heading in transactions that span several chapters; 77 percent of these lines belong to transactions with 11 or more products.

[10]Analogously, we measure exporter size by the annual value of a firm's export shipments.

of importers with at least one MPT rises from 58 percent in the bottom quartile of the size distribution to 91 percent in the top quartile.

Figure 6: Firm Size and Bundling Shares



Notes: The sample is importers with at least one MPT in the year, 2011–2020 (362,198 importer-years). The panels plot the MPT share of the importer's transactions, import value, and quantity in the year against importer size, defined as in Figure 5. Dots are means within 100 equal-count bins of importer size. The line is the fit of a fractional logit of the share on log importer size, estimated on the unbinned importer-years and drawn between the 0.5th and 99.5th percentiles of importer size. The quantity panel plots the within-product MPT share of quantity, computed as in Table 1; it leaves out seven importer-years without a usable quantity.

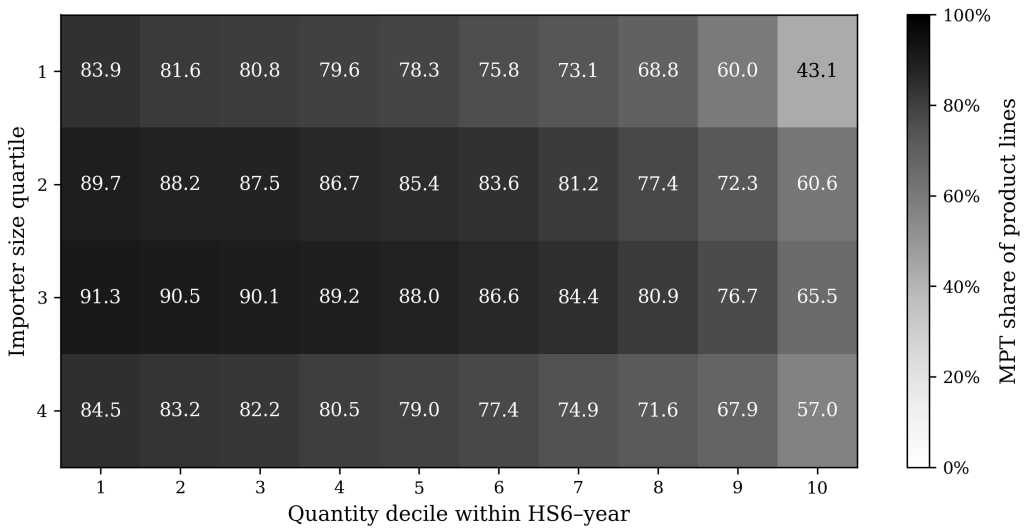
Yet conditional on conducting at least one MPT in the year, the share of a firm's transactions that are multiproduct declines sharply with firm size. Figure 6 shows that the MPT share of transactions falls from 87 percent in the smallest of 100 size bins to 27 percent in the largest. In the same size quartiles as above, MPTs account on average for 73 percent of these importers' transactions in the bottom quartile and for 41 percent in the top quartile. The MPT shares of import value and of quantity decline in the same way, from 77 to 46 percent. Thus, while large firms are more likely to engage in MPTs at least once, among the firms that do, small firms rely much more heavily on them.

Fact 3: Conditional on firm size, larger purchases of a product are less likely to be in an MPT.

The variation in MPT intensity across firm size could in principle reflect differences in purchase size. Large firms tend to make large purchases, which might make it worthwhile to search for the lowest-cost supplier of each product. Figure 7 disaggregates MPT frequency by both firm and purchase size. The unit of observation is the product line, an HS6 product within a transaction. We rank each line by its quantity among all lines of the same HS6 product and tariff unit in the same year and report the shares of lines that belong to an MPT by decile of this ranking and by quartile of importer size.[11] Within each size quartile, the likelihood that a product is purchased in an MPT declines monotonically with the size of the purchase. Among importers in the bottom size quartile, the MPT share of product lines

[11] We measure purchase size for the product line rather than for the transaction because the value of an MPT is the sum of several product lines by definition, so the size of a transaction and the size of the purchases in it are different objects. Ranked by total transaction value (Figure A1 in the Appendix A), the MPT share of transactions is lowest in the bottom decile in every size quartile (17 to 26 percent), rises through the fifth decile, and is between 34 and 51 percent in the top two deciles. In every decile it is lowest for importers in the top size quartile.

Figure 7: MPT Frequency by Firm and Purchase Size



Notes: Each cell reports the percentage of product lines that belong to an MPT, in the sample of Section 2, 2011–2020. A product line is an HS6 product within a transaction; all HS6 products are included, and quantities are compared only within a product and tariff unit. Rows are quartiles of importer size, defined as in Figure 5, with cutpoints computed over all importer-years. Columns are deciles of the line's quantity among all lines of the same HS6 product and tariff unit in the same year; lines with equal quantities share a decile, so deciles need not contain equal numbers of lines. HS6–unit–year cells with fewer than 10 lines (0.01 percent of lines with a usable quantity) are excluded, and lines without a usable quantity (1.0 percent) are left out.

falls from 86 percent in the lowest quantity decile to 52 percent in the highest. Among the largest importers the same pattern is evident, the share falls from 88 to 50 percent.

Figure 7 holds fixed only the product, the year, and the importer's size quartile. Table 4 makes the comparison holding the importer-year and the exporter-product-year fixed (Fact 1). Across size quartiles, by contrast, the MPT share of product lines is not monotone. Within every quantity decile it rises from the bottom quartile to the third and is lower in the top quartile than in the second and third. This ordering differs from Fact 2, which is restricted to importers with at least one MPT in the year, counts transactions rather than product lines, and weights importers equally.

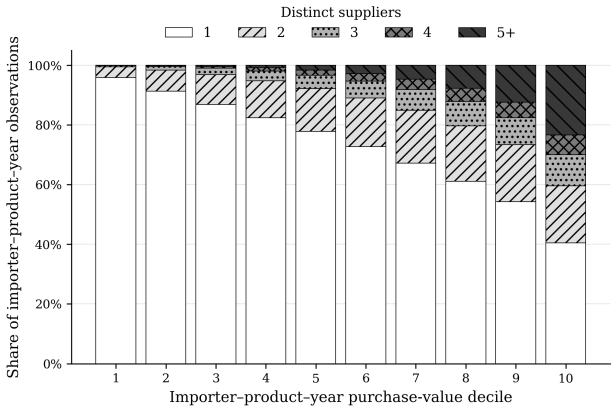
Fact 4: The number of suppliers per product rises with firm size.

Figure 8a shows the number of suppliers from which an importer sources a given product in a given year, across the distribution of purchase sizes. Purchase size is here the value of the importer's purchases of the product in the year, summed over its transactions. The smallest annual purchases almost always come from a single supplier, 96 percent of those in the bottom decile do. Among annual purchases in the top decile, 40 percent come from a single suppliers, 19 percent from two, and 40 percent from three or more. In the top decile, a majority of annual purchases thus come from two or more exporters.

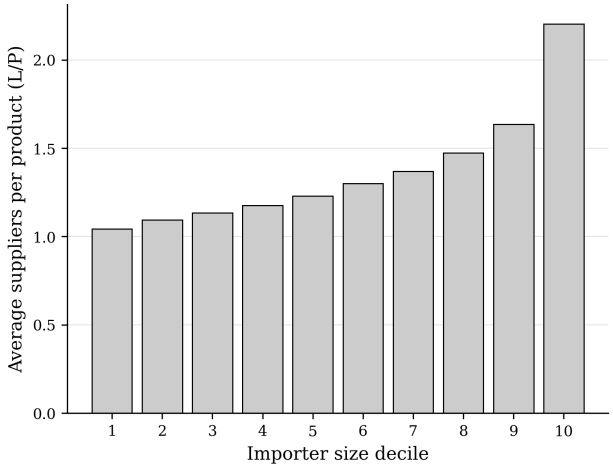
Figure 8b shows that the average number of suppliers per product rises monotonically with firm size, from 1.04 in the bottom size decile to 2.20 in the top decile. In each of the bottom five deciles, the median importer has exactly one supplier for each of its products. In the top decile, the median importer has 1.67 suppliers per product. At the same time, importers in every size decile have fewer

Figure 8: Suppliers per Product and Products per Supplier (Mexican Importers)

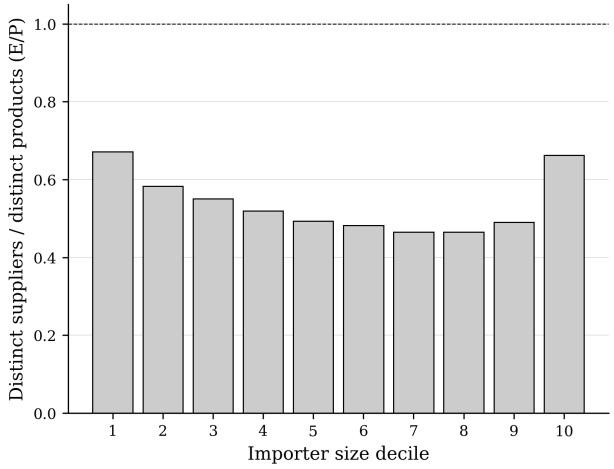
(a) Suppliers per Product, by Importer–Product Size



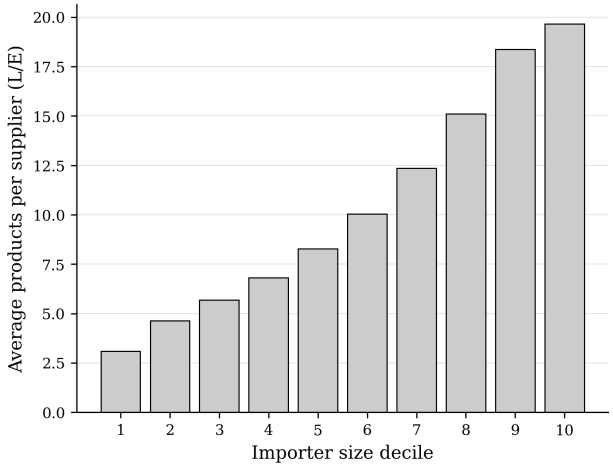
(b) Suppliers per Product (L/P), by Importer Size



(c) Suppliers to Products (E/P), by Importer Size



(d) Products per Supplier (L/E), by Importer Size



Notes: Mexican import transactions in the sample of Section 2, 2011–2020. Suppliers and products are counted within a calendar year. In panel (a), each observation is an HS6 product imported by a firm in a given year (14,129,119 observations). Observations are sorted into deciles of the value of the importer's purchases of the product in the year, with cutpoints computed over all observations, and the bars show the shares of observations with one, two, three, four, and five or more distinct suppliers. In panels (b) to (d), each observation is an importer in a given year (471,504 importer-years), sorted into deciles of the importer's annual import value as defined in Figure 5, in US dollars rather than relative to the annual mean, with cutpoints computed over all importer-years. For each importer-year, P is the number of distinct HS6 products, E the number of distinct suppliers, and L the number of distinct supplier–product pairs. The panels plot the means across importer-years of L/P, the average number of suppliers per product; E/P, the ratio of suppliers to products (the dashed line marks one); and L/E, the average number of products per supplier.

suppliers than products on average. Figure 8c plots the ratio of an importer's number of distinct suppliers to its number of distinct products. The average ratio is 0.67 in the bottom decile, falls to 0.46 in the seventh decile, and is 0.66 in the top decile. In the top decile the median ratio is 0.35, well below the mean, and the share of importers with fewer suppliers than products rises in every decile, from 53 percent in the bottom decile to 86 percent in the top. Each supplier provides more than one product on average in every size decile, and more to larger importers. The average number of products that an importer buys from each of its suppliers rises from 3.1 in the bottom decile to 19.7 in the top decile (Figure 8d).[12]

4 Model

4.1 Demand

Consider a global sourcing environment in which domestic importing firms demand two intermediate inputs, 1 and 2, that can be sourced from global suppliers. An importing firm's demand is represented by the random vector $\mathbf{q} = (q_1, q_2)$ with density $f(\mathbf{q}, \lambda)$, where $\lambda \geq 1$ denotes firm size (or underlying productivity). We assume demands are homothetic in firm size, taking the form $f(\mathbf{q}, \lambda) = \lambda^{-2} g(\mathbf{q}/\lambda)$, with support $q_i \in [\lambda \underline{q}_i, \lambda \bar{q}_i]$ and $g(\cdot)$ a baseline density. This scaling structure implies that larger firms are simply scaled-up versions of smaller firms: expected demand for each input grows in direct proportion to firm size, $E(q_i|\lambda) = \lambda E(q_i|1)$. An importer incurs a fixed sourcing cost F for each supplier it visits.

4.2 Supply

There are two potential suppliers and two inputs, each indexed by $i, j \in \{1, 2\}$. Let p_{ij} denote the price charged by supplier i for input j. Supplier i charges a lower price for input i, which we refer to its "core" product. Thus, for $j \neq i$, $p_{ii} < p_{ji}$, $\Delta_i \equiv p_{ji} - p_{ii} > 0$. The price gap Δ_i measures supplier i's per-unit price advantage in its core input, or equivalently, the price premium paid when input i is purchased from the non-core supplier j.

4.3 Multiproduct Sourcing

Because the fixed sourcing cost F does not vary with the quantity purchased, the importer faces a trade-off between sourcing each input from its cheapest supplier and avoiding the additional cost of visiting multiple suppliers. If the demand for each input is sufficiently large, the savings from accessing its lower-cost supplier outweigh the fixed cost, and the importer will choose to split purchases across suppliers. Otherwise, it is optimal to source both goods from a single supplier, even if this means paying a higher per-unit price for one of the inputs.

[12] These statistics count an importer's suppliers and products within a calendar year. Counted over 2011-2020 instead (Figure A2 in Appendix A), the number of suppliers per product rises from 1.09 in the bottom size decile to 2.97 in the top decile; the ratio of suppliers to products is 0.56 in the bottom decile, 0.40 in the eighth, and 0.52 in the top decile; and 59 percent of the importer-product pairs in the top decile of purchase value have three or more suppliers.

Formally, fix an input $j \in \{1, 2\}$, and let $i \neq j$ denote the other supplier. Because input j is supplier j's core product, purchasing it from supplier i rather than supplier j entails the per-unit price premium

$$p_{ij} - p_{jj} = \Delta_j > 0.$$

Suppose that the importer already visits supplier i. Adding supplier j solely to purchase input j would reduce variable expenditure by $\Delta_j q_j$, but would require the importer to incur the additional fixed sourcing cost F. The importer therefore chooses not to visit supplier j when

$$\Delta_j q_j < F \quad \Longleftrightarrow \quad q_j < q_j^0 \equiv \frac{F}{\Delta_j}.$$

In this case, the savings from purchasing input j from its core supplier do not cover the fixed cost of adding that supplier. It is therefore cheaper to purchase input j from supplier i, despite its higher per-unit price. The threshold q_j^0 rises with the fixed sourcing cost F and falls with the price advantage Δ_j: a higher fixed cost makes adding supplier j less attractive, whereas a larger price advantage makes doing so more attractive.

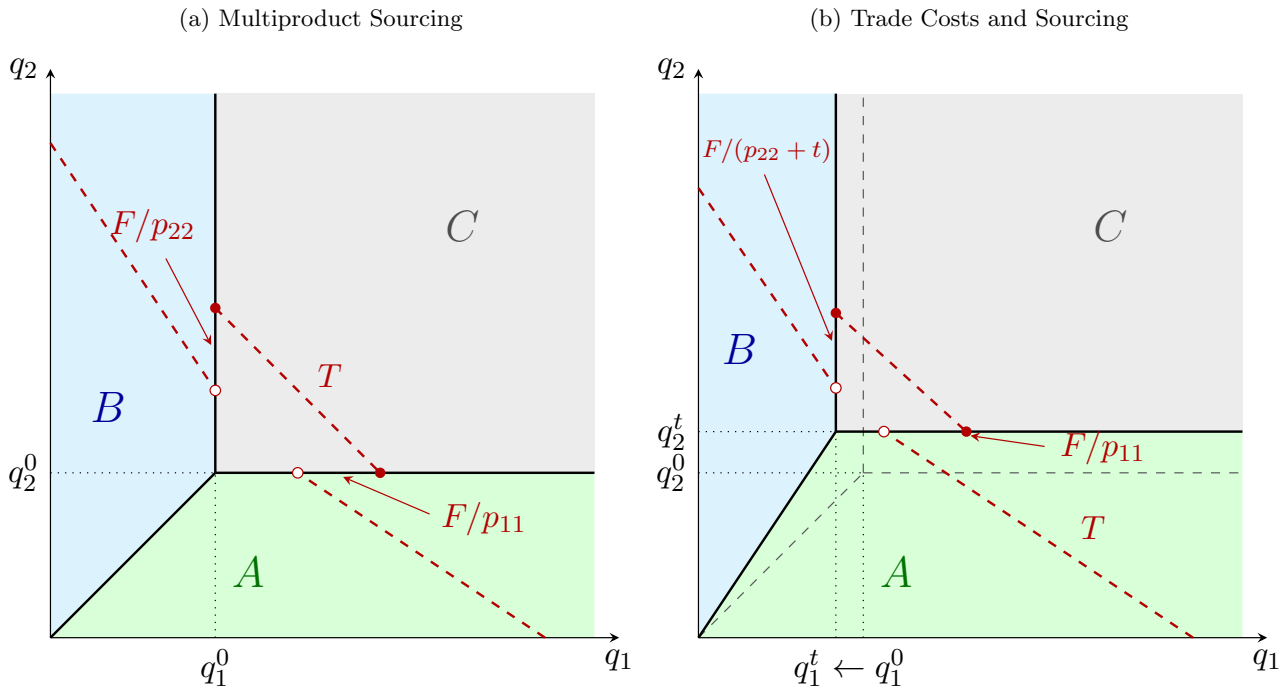
If both demands are below their respective thresholds, $q_1 < q_1^0$ and $q_2 < q_2^0$, the importer can either source both inputs from supplier 1 or both from supplier 2. The choice depends on relative price penalties: sourcing both from supplier 1 is optimal when

$$\Delta_2 q_2 \leq \Delta_1 q_1 \quad \Leftrightarrow \quad q_1 \geq \frac{\Delta_2}{\Delta_1} q_2,$$

while otherwise the importer prefers supplier 2.[13]

[13] We assume that the values of $\underline{q}_i$ are sufficiently large so that the surplus from visiting a single supplier is positive at $(\underline{q}_1, \underline{q}_2)$ when $\lambda = 1$. This ensures that purchases take place for all λ and realizations of $\mathbf{q}$.

Figure 9: Transaction Size and Sourcing Strategy



Notes: Panel (a) depicts the importer's sourcing strategy across demand realizations (q_1, q_2). In region A (B), the importer purchases both inputs from supplier 1 (2) in a multiproduct transaction: demand for input 2 (1) is below the threshold q_2^0 (q_1^0), and supplier 1 (2) is the lower-cost single source. In region C, demand for each input is at or above its threshold, and the importer sources each input from its lowest-price supplier in two single-product transactions. The red dashed segments trace an iso-expenditure locus: combinations of (q_1, q_2) with the same total expenditure T, excluding fixed sourcing costs. The slope of each segment is minus the relative price paid for input 1, which differs across regions because the importer pays a premium on a supplier's non-core input. The locus jumps by $\frac{F}{p_{11}}$ and $\frac{F}{p_{22}}$ at the boundaries of region C, where adding a second supplier lowers expenditure by F. Panel (b) shows an increase in t, the per-unit trade cost of supplier 2 relative to supplier 1, drawn with the trade cost of supplier 1 set to zero. Gray dashed lines denote the boundaries at $t = 0$; solid lines and shading denote the boundaries and regions at $t > 0$. The thresholds move from (q_1^0, q_2^0) to (q_1^t, q_2^t), and the boundary between regions A and B steepens.

Figure 9a illustrates the importer's sourcing strategy across demand realizaitons, where the horizontal axis measures demand for input 1 (q_1) and the vertical axis measures demand for input 2 (q_2). Three boundaries divide the space: the thresholds q_1^0 and q_2^0, and the ray $\Delta_2 q_2 = \Delta_1 q_1$ along which the importer is indifferent between the two suppliers as a single source. Because $\Delta_1 q_1^0 = \Delta_2 q_2^0 = F$, this ray passes through (q_1^0, q_2^0), where the importer is indifferent among all three sourcing options.

In region A (green), the importer purchases both inputs from supplier 1 in a multiproduct transaction. Demand for input 2 is too small to justify visiting supplier 2, $q_2 < q_2^0$, and consolidating with supplier 1 entails no greater expenditure penalty than consolidating with supplier 2, $\Delta_2 q_2 \leq \Delta_1 q_1$. The latter condition holds automatically when $q_1 \geq q_1^0$, since $\Delta_1 q_1 \geq F > \Delta_2 q_2$. Region B (blue) follows the same logic with suppliers reversed: the importer purchases both inputs from supplier 2 when $q_1 < q_1^0$ and $\Delta_1 q_1 < \Delta_2 q_2$. In region C (gray), demand for each input is large enough to justify sourcing it from its lowest-price supplier, $q_1 \geq q_1^0$ and $q_2 \geq q_2^0$. The importer therefore visits both suppliers and conducts two single-product transactions.

The red dashed lines in Figure 9a trace an iso-expenditure locus: the combinations of (q_1, q_2) at

which the importer's expenditure on the two inputs, excluding fixed sourcing costs, equals T. In regions A and B, T is the value of the single multiproduct transaction. In region C, it is the combined value of the two single-product transactions. The slope of each segment equals minus the relative price that the importer pays for input 1: $-\frac{p_{11}}{p_{12}}$ in region A, $-\frac{p_{11}}{p_{22}}$ in region C, and $-\frac{p_{21}}{p_{22}}$ in region B. The change in slope across regions reflects the fact that in a multiproduct transaction the importer pays a higher price for the supplier's non-core input. The locus is also discontinuous at the boundaries of region C. Along the boundary between regions A and C, where $q_2 = q_2^0$, the importer is indifferent between a multiproduct transaction with supplier 1 and two single-product transactions: adding supplier 2 lowers expenditure by $\Delta_2 q_2^0 = F$, which excatly offsets the additional fixed cost. For a given expenditure T, the locus therefore shifts right by $\frac{F}{p_{11}}$ at this boundary. By the same argument, it shifts up by $\frac{F}{p_{22}}$ at the boundary between regions B and C, where $q_1 = q_1^0$. The figure draws the locus for $T > p_{11}q_1^0 + p_{12}q_2^0$, in which case its segment in region A ends at $q_2 = q_2^0$; for lower values of T, that segment ends at the ray.

Using the threshold levels and the homothetic scaling of demand, we obtain

Proposition 1 *The probability that an importer of size λ engages in a multiproduct transaction is non-increasing in λ. It is strictly decreasing whenever $q_j^0/\lambda \in (\underline{q}_j, \bar{q}_j)$ for some $j \in \{1,2\}$ and g is strictly positive on its support.*

Proof. Using a change of variable, the probability that the importer visits both suppliers is

$$\Pr(q_1 \geq q_1^0 \wedge q_2 \geq q_2^0 | \lambda) \quad = \quad \int_{q_1^0}^{\lambda \bar{q}_1} \int_{q_2^0}^{\lambda \bar{q}_2} \lambda^{-2} g(\mathbf{q}/\lambda) dq_2 dq_1 = \int_{q_1^0/\lambda}^{\bar{q}_1} \int_{q_2^0/\lambda}^{\bar{q}_2} g(\mathbf{q}') dq_2' dq_1' \tag{1}$$

which is non-decreasing in λ because the lower limits of integration falls as λ rises. The probablity of a multiproduct transaction is the complement of this probability. ■

Remark 4.1 *This result is consistent with the decline in MPT intensity across firm size documented in Fact 2.*

The negative relationship between firm size and the likelihood of multiproduct transactions is not specific to the homothetic scaling assumption above. In fact, it would also arise under an alternative specification such as $f(\mathbf{q}, \lambda) = g(q_1 - \lambda, q_2 - \lambda)$, where demand for each input increases additively with λ. In this case, the support of the distribution shifts to the right linearly with firm size, ensuring that larger firms are still more likely to exceed the thresholds for sourcing from multiple suppliers.

The second result follows immediately from the properties of the multiproduct and single product regions in Figure 9a.

Proposition 2 *(i) In a multiproduct transaction, the importer pays more for the supplier's non-core product than it would if that product were purchased from the supplier for which it is the core good. (ii) The quantity of the non-core product in multiproduct transactions is lower than it would need to be for the importer to find it optimal to engage in single-product transactions with both suppliers.*

4.3.1 Symmetric Sourcing

To characterize how the share of multiproduct transactions varies with transaction value, we derive the densities of single product and multiproduct transactions of a given value T for an importer size λ. For tractability, we focus on the symmetric case where $\Delta_1 = \Delta_2$ and the demand distribution is exchangeable, $g(x, y) = g(y, x)$.

Let $h_1(T, \lambda)$ denote the density of single product transactions (SPTs) of total value T between an importer of size λ and supplier 1. In such a transaction, the importer purchases input 1 (the supplier's core good) in quantity $q_1 = T/p_{11}$, while the demand for input 2 must be sufficiently large to justify visiting supplier 2, i.e. $q_2 \geq q_2^0$. Accordingly,

$$h_1(T, \lambda) = \frac{1}{p_{11}} \int_{q_2^0}^{\lambda \bar{q}_2} f\left(\frac{T}{p_{11}}, q_2, \lambda\right) dq_2 = \frac{1}{\lambda p_{11}} \int_{q_2^0/\lambda}^{\bar{q}_2} g\left(\frac{T}{\lambda p_{11}}, q_2'\right) dq_2'. \tag{2}$$

when $\max\{\lambda \underline{q}_1, q_1^0\} \leq \frac{T}{p_{11}} \leq \lambda \bar{q}_1$. An analogous expression holds for SPTs with supplier 2.

We now turn to multiproduct transactions (MPTs), where the importer purchases both inputs from the same supplier. In this case, the total transaction value T includes expenditures on both goods rather than a single input. For an MPT with supplier 1, the total value is given by $T = p_{11}q_1 + p_{12}q_2$. For relatively small transactions satisfying $T \leq p_{11}q_1^0 + p_{12}q_2^0$, at least one input demand falls below its threshold, making it suboptimal for the importer to visit both suppliers. In this region, the importer therefore consolidates purchases with a single supplier. The choice between suppliers depends on relative price penalties. Sourcing both goods from supplier 1 is optimal when $\Delta_2 q_2 \leq \Delta_1 q_1$, meaning that the premium paid on good 2 (the non-core product) at supplier 1 does not exceed the premium that would be paid on good 1 at supplier 2. The upper bound of q_2 consistent with supplier 1 being chosen as the single supplier follows from the indifference condition $\Delta_2 q_2 = \Delta_1 q_1$. Solving this condition and $T = p_{11}q_1 + p_{12}q_2$ jointly yields $q_2^{\max} = \frac{T\,\Delta_1}{\Delta_2 p_{11} + \Delta_1 p_{12}}$, where the denominator aggregates the non-core price penalties associated with each supplier.

For large transactions, $T > p_{11}q_1^0 + p_{12}q_2^0$, at least one of the two input demands exceeds its threshold, so the importer may purchase from both suppliers (two SPTs). Nevertheless, a multiproduct purchase from supplier 1 still arises when demand for good 2 remains below q_2^0, in which case the importer chooses to buy both inputs from supplier 1 despite the higher price on good 2. The upper bound of q_2 is then q_2^0, as in Figure 9a, where the segment of the locus in region A ends at $q_2 = q_2^0$.

Formally, let $h_{11}(T, \lambda)$ denote the density of multiproduct transactions in which an importer of size λ purchases both goods from supplier 1 and the total transaction value equals T. For any candidate value of q_2, the corresponding q_1 is implied by the transaction identity $q_1 = (T - p_{12}q_2)/p_{11}$. Accordingly,

$$h_{11}(T, \lambda) = \begin{cases} \dfrac{1}{p_{11}} \displaystyle\int_{\lambda \underline{q}_2}^{\frac{T\Delta_1}{\Delta_2 p_{11} + \Delta_1 p_{12}}} f\left(\frac{T - p_{12}q_2}{p_{11}}, q_2, \lambda\right) dq_2, & \text{if } T \leq p_{11}q_1^0 + p_{12}q_2^0, \\ \dfrac{1}{p_{11}} \displaystyle\int_{\lambda \underline{q}_2}^{q_2^0} f\left(\frac{T - p_{12}q_2}{p_{11}}, q_2, \lambda\right) dq_2, & \text{if } T > p_{11}q_1^0 + p_{12}q_2^0. \end{cases} \tag{3}$$

Under the symmetry assumption, the share of supplier 1's transactions that are multiproduct can be expressed as

$$\mu_1(T,\lambda) = \frac{h_{11}(T,\lambda)}{h_{11}(T,\lambda) + h_1(T,\lambda)}.$$

This ratio represents the conditional probability that a transaction involving supplier 1 is multiproduct, given importer size λ and transaction value T. Although symmetry is a simplifying case, it provides sufficient tractability to characterize how the share of MPTs varies systematically with firm size (λ) and transaction size (T). These expressions also highlight that the firm-size and transaction-size results are conceptually distinct. Proposition 1 establishes that the unconditional probability of an MPT is weakly decreasing in firm size λ: as demands increase relative to the fixed thresholds, both inputs are more likely to be sourced separately. By contrast, the relationship between transaction value T and the conditional MPT share $\mu_1(T,\lambda)$ depends on the correlation between input demands, as the following cases illustrate.

Demand correlation. Consider first the case in which the demands for the two goods are perfectly positively correlated, $q_1 = q_2 = q$. In this setting, total transaction value for an MPT with supplier 1 is $T = (p_{11} + p_{12})q$, and the implied distribution of MPTs can be written as $\phi(q,\lambda) = f(q,q,\lambda)$. When q is small—so that both input demands are below their respective thresholds—the importer consolidates purchases into a single multiproduct transaction and $\mu_1(T,\lambda) \approx 1$. As q (and equivalently T) increases, both input demands rise proportionally and eventually exceed their thresholds, leading the importer to split purchases across suppliers. Consequently, the share of MPTs, $\mu_1(T,\lambda)$, declines with transaction size. This qualitative pattern aligns with Fact 3: in the data, larger transactions are less likely to be multiproduct.

At the other extreme, suppose demands are perfectly negatively correlated, $q_1 + q_2 = k(\lambda)$. When demand for one input is high, demand for the other is low, implying that both inputs rarely exceed their thresholds simultaneously. In this case, importers frequently find it optimal to consolidate purchases with a single supplier, and the resulting MPTs tend to have higher total values than comparable single-product transactions, since they include at least one high-demand input. More generally, when demand shocks across inputs are positively correlated, firms with larger demand realizations are more likely to purchase goods in separate transactions, whereas with negatively correlated shocks, consolidation into multiproduct transactions remains more prevalent.

4.4 Trade Costs and MPTs

To examine the role of differential trade costs for multiproduct transactions, let τ_i denote the per-unit trade (delivery) cost associated with supplier i, so that the delivered price of input j from supplier i is $p_{ij} + \tau_i$, and define $t \equiv \tau_2 - \tau_1$ as the trade cost of supplier 2 relative to supplier 1 (so $t > 0$ indicates that supplier 2 is relatively more costly). Trade costs enter the sourcing problem only through the delivered price premia on non-core inputs:

$$\Delta_1 \;=\; p_{21} - p_{11} + t, \qquad \Delta_2 \;=\; p_{12} - p_{22} - t,$$

where Δ_i denotes the per-unit penalty of purchasing input i from the non-core supplier.[14] The threshold quantities below which the importer does not find it profitable to visit an additional supplier are therefore

$$q_1^t = \frac{F}{\Delta_1} = \frac{F}{(p_{21} - p_{11}) + t}, \qquad q_2^t = \frac{F}{\Delta_2} = \frac{F}{(p_{12} - p_{22}) - t}. \tag{4}$$

which equal q_1^0 and q_2^0 at $t = 0$. Hence an increase in t (i.e. supplier 2 becoming relatively more costly) raises the threshold q_2^t and lowers q_1^t: the importer requires a larger demand for good 2 to justify visiting supplier 2, while the range of demands for which supplier 1 is the single-supplier optimal choice expands. It also steepens the ray $\Delta_2 q_2 = \Delta_1 q_1$, whose slope Δ_1/Δ_2 rises, so that supplier 1 is the preferred single source for a wider set of demands.

In Figure 9b, the gray dashed lines represent the sourcing boundaries when $t = 0$, and the solid lines correspond to the new boundaries after an increase in the relative trade cost of supplier 2 ($t > 0$). The figure is drawn for $\tau_1 = 0$ and $\tau_2 = t$, and the shaded regions A, B, and C are those at $t > 0$. The thresholds move from (q_1^0, q_2^0) to (q_1^t, q_2^t), and the boundary between regions A and B steepens. The red dashed segments again trace an iso-expenditure locus, now valued at delivered prices; its discontinuities at the boundaries of region C are $\frac{F}{p_{11}}$ and $\frac{F}{(p_{22}+t)}$. Letting $\mathbf{q} = (q_1, q_2)$ denote the vector of input demands, the conditional probability that a transaction involving supplier 1 is a multiproduct transaction is

$$\frac{\Pr(\mathbf{q} \in A)}{\Pr(\mathbf{q} \in A \cup C)} = \frac{\Pr(\mathbf{q} \in A)}{1 - \Pr(\mathbf{q} \in B)},$$

and similarly, the conditional probability that a transaction involving supplier 2 is a multiproduct transaction is

$$\frac{\Pr(\mathbf{q} \in B)}{\Pr(\mathbf{q} \in B \cup C)} = \frac{\Pr(\mathbf{q} \in B)}{1 - \Pr(\mathbf{q} \in A)}.$$

An increase in t makes supplier 2 relatively more costly, which unambiguously raises $\Pr(\mathbf{q} \in A)$ (region A expands) and lowers $\Pr(\mathbf{q} \in B)$ (region B contracts). The effect on $\Pr(\mathbf{q} \in C)$, corresponding to two single-product transactions, is ambiguous because the relative magnitudes of the expansion of region A and the contraction of region B determine whether region C expands or contracts.

Let $x = \Pr(\mathbf{q} \in A)$ and $y = \Pr(\mathbf{q} \in B)$, with $x + y < 1$. The share of supplier 1's transactions that are multiproduct is then $\mu_1 = x/(1 - y)$. Differentiating with respect to t gives

$$d\mu_1 = \frac{(1 - y)\, dx + x\, dy}{(1 - y)^2},$$

which is positive when

$$dx > -\frac{x}{1 - y}\, dy. \tag{5}$$

Since an increase in t implies $dx > 0 > dy$, inequality (5) requires that the decline in the frequency of multiproduct transactions for supplier 2 not be too large relative to the expansion of multiproduct transactions for supplier 1.

[14] We restrict t to $-(p_{21} - p_{11}) < t < p_{12} - p_{22}$, so that both premia remain positive and each supplier retains a cost advantage in its core input.

Analogously, the share of supplier 2's transactions that are multiproduct, $\mu_2 = y/(1-x)$, falls when

$$dx < -\frac{1-x}{y}\, dy. \tag{6}$$

For both effects to hold simultaneously—that is, for μ_1 to rise and μ_2 to fall—the rate of change must satisfy

$$-\frac{dx}{dy} \in \left(\frac{x}{1-y}, \frac{1-x}{y}\right). \tag{7}$$

We are also interested in how trade costs shape the *difference* in the prevalence of multiproduct transactions between the two suppliers, defined as $\Delta\mu = \mu_1 - \mu_2$. This difference increases ($d\Delta\mu > 0$) if and only if

$$\left(\frac{1}{1-y} - \frac{y}{(1-x)^2}\right) dx - \left(\frac{1}{1-x} - \frac{x}{(1-y)^2}\right) dy > 0. \tag{8}$$

Starting from a symmetric initial position, $x = y = s$, we obtain $d\Delta\mu = \frac{1-2s}{(1-s)^2}(dx - dy)$. Because $x + y < 1$ implies $s < 1/2$, and an increase in t yields $dx > 0 > dy$, it follows that $d\Delta\mu > 0$. Thus, starting from symmetry, an increase in supplier 2's relative trade cost increases the MPT-share difference in favor of supplier 1.

4.5 Strategic Pricing and Exporter Behavior

Section 4.2 assumes that suppliers charge fixed prices for their core and peripheral products. However, exporters may strategically adjust prices in response to importer demand to maximize profits subject to deviation constraints that prevent importers from purchasing from alternative suppliers. This section incorporates strategic pricing into the benchmark model.

As above, each importer incurs a fixed cost of F to visit an exporter. Upon a visit, the exporter observes demand and chooses prices strategically. The importer has the option of buying both products from the exporter, buying one product from the exporter and visiting the other exporter or buying neither product from the exporter. The exporter's problem is to choose prices to maximize profits subject to the deviation constraints of the importer. We assume throughout that the importer knows which product is the exporter's core product prior to the visit.

4.5.1 Importer Demand and Surplus

Each importer has demand for two inputs, with the value of product j to the importer is given by $A_j q_j^\theta$ for $j \in \{1, 2\}$ and $\theta \in (0, 1)$. The parameter θ is a demand elasticity parameter that governs the diminishing marginal value of the product to the importer. Lower values of θ imply more inelastic demand. The vector $\mathbf{A}$ is a random vector with pdf $f(\mathbf{A}, \lambda) = \mu^{-2} g(\mathbf{A})$, where μ reflects the size of the firm.

The importer's surplus from purchasing q_j units of good j, net of its expenditure $p_j q_j$, is

$$A_j q_j^\theta - p_j q_j$$

which yields demand, $d_j(p_j) = \left(\frac{\theta A_j}{p_j}\right)^{\frac{1}{1-\theta}}$, and realized surplus, $S_j(A_j, p_j) = \kappa A_j^{\frac{1}{1-\theta}} p_j^{\frac{\theta}{\theta-1}}$, where $\kappa = \theta^{\frac{\theta}{1-\theta}} - \theta^{\frac{1}{1-\theta}} > 0$ is a constant that depends only on the elasticity parameter θ.

4.5.2 Deviation Constraints and Equilibrium Pricing

As before, we consider two exporters with constant marginal costs of producing two products. Product i is the core product of exporter i: exporter i produces product i at marginal cost c and product $j \neq i$ at marginal cost Δc, where $\Delta > 1$. Exporter 1's profit is

$$\pi_1(p_{11}, p_{12}) = (p_{11} - c)d_1(p_{11}) + (p_{12} - \Delta c)d_2(p_{12}).$$

Absent any deviation constraints, exporter 1 charges the unconstrained monopoly price for each product:

$$p_{11}^m = \frac{c}{\theta}, \qquad p_{12}^m = \frac{\Delta c}{\theta}.$$

Suppose an importer visits exporter 1 and considers purchasing both products from it. To sustain this multiproduct transaction, exporter 1 must ensure that the importer does not prefer either to add exporter 2 solely for good 2 or to switch both purchases to exporter 2.

Single-product deviation constraint (adding exporter 2 for good 2 only)

Under the first deviation, the importer continues to purchase good 1 from exporter 1 but visits exporter 2 to purchase good 2. Because good 1 is purchased from exporter 1 under either arrangement, its surplus cancels from the comparison. Purchasing good 2 from exporter 2 raises the importer's surplus from $S(A_2, p_{12})$ to $S(A_2, p_{22}^e)$, but requires paying the additional sourcing cost F. The no-deviation constraint is therefore

$$S(A_2, p_{12}) - S(A_2, p_{22}^e) + F \geq 0, \tag{9}$$

where p_{22}^e is the price that the importer expects exporter 2 to charge for its core product, good 2. Equivalently, the surplus gain from purchasing good 2 from exporter 2 cannot exceed the cost of adding that exporter. Constraint (9) therefore places an upper bound on p_{12}. When it binds, exporter 1 charges the highest price that retains the sale of good 2 within the multiproduct transaction.[15]

[15] Using the expression for realized surplus derived above, constraint (9) can be written as

$$\kappa A_2^{\frac{1}{1-\theta}} \left[p_{12}^{\frac{\theta}{\theta-1}} - (p_{22}^e)^{\frac{\theta}{\theta-1}} \right] + F \geq 0.$$

Whenever

$$(p_{22}^e)^{\frac{\theta}{\theta-1}} - \frac{F}{\kappa A_2^{\frac{1}{1-\theta}}} > 0,$$

this constraint places the following upper bound on exporter 1's non-core price:

$$p_{12} \leq \bar{p}_{12}(A_2) \equiv \left[(p_{22}^e)^{\frac{\theta}{\theta-1}} - \frac{F}{\kappa A_2^{\frac{1}{1-\theta}}} \right]^{\frac{\theta-1}{\theta}}. \tag{10}$$

Thus, $\bar{p}_{12}(A_2)$ is the highest price at which exporter 1 can retain the sale of good 2. If the expression inside the brackets is nonpositive, the deviation threat does not impose a finite upper bound, and exporter 1 charges its unconstrained

Multiproduct deviation constraint (switching both goods to exporter 2)
Alternatively, an importer should not find it profitable to purchase both goods from exporter 2:

$$S(A_1, p_{11}) + S(A_2, p_{12}) - S(A_1, p_{21}^e) - S(A_2, p_{22}^e) + F \geq 0 \tag{11}$$

This constraint prevents the importer from defecting to exporter 2 entirely to purchase both goods.

4.5.3 Demand-Dependent Pricing

Equilibrium prices depend on the importer's demand parameters (A_1, A_2) whose distribution varies with firm size λ. Exporter 1 will optimally price its peripheral product at different levels depending on the importer's demand.

Suppose that the expectation of an importer who is currently visiting exporter 1 is that $p_{22}^e = \frac{c}{\theta}$. The exporter will strategically choose one of three pricing options.

1. **Low-demand importers** $(A_2 \leq \tilde{A}_2)$: For importers with sufficiently low demand for the peripheral product, the importer's switching threat is weak—the surplus it would obtain by visiting exporter 2 is small. In this case, exporter 1 sells both goods to the importer at the unconstrained price vector

$$\left(p_{11} = \frac{c}{\theta},\ p_{12} = \frac{\Delta c}{\theta}\right) \text{ for } A_2 \leq \tilde{A}_2$$

2. **Intermediate-demand importers** $(\tilde{A}_2 < A_2 < \hat{A}_2)$: For importers with intermediate demand, the exporter offers a discount on the peripheral product relative to the monopoly level, pricing to just keep the importer from deviating. Exporter 1 can profitably deter deviation to exporter 2 by choosing a price such that (9) is satisfied with equality as long as $p_{12} > \Delta c$. There are two cases to consider, depending on whether or not the inefficiency in producing the peripheral good exceeds the monopoly markup.

   (i) If $\Delta c \leq \frac{c}{\theta}$, then there exists a threshold price for each importer that will deter deviation to exporter 2 for good 2 with that threshold approaching $\frac{c}{\theta}$ as $A_2 \to \infty$.

   (ii) If $\Delta c > \frac{c}{\theta}$, then there will exist a threshold value $\hat{A}_2$ such that (9) is satisfied with equality at $p_{12} = \Delta c$.

   Ceteris paribus, higher fixed sourcing costs F strengthen the importer's incentive to consolidate, allowing the exporter to charge a higher price. For a given θ, higher peripheral demand A_2 makes deviation more attractive and therefore requires a deeper discount to retain the importer.

3. **High-demand importers** $(A_2 > \hat{A}_2)$ If $\Delta c > \frac{c}{\theta}$ then among importers with sufficiently high demand for the peripheral product, exporter 1 cannot profitably offer a sufficient discount to

monopoly price.

retain the importer. In this case, the importer will purchase each good in separate transactions from different exporters.

Allowing the exporter to discount the peripheral good to importers with a larger demand raises the SPT threshold, but still provides a threshold independent of core product demand so that larger importers engage in more SPTs.

4.5.4 Minimum Sourcing Thresholds

We have established the payoffs to an importer from visiting exporter 1 under the assumption that $p^e_{22} = \frac{c}{\theta}$ and shown that the optimal price for good 1 will be $p_{11} = \frac{c}{\theta}$. It then follows by the symmetry of the exporters that our assumption on exporter 2's price is optimal for exporter 2 given the expectation that $p_{11} = \frac{c}{\theta}$. Since an importer pays a fixed cost of visiting an exporter, it remains to establish the threshold demand levels for an importer to visit an exporter.

Because the marginal importer will be making a multiproduct purchase, the price of the exporter's core good will be $\frac{c}{\theta}$ and the price of the other good will be $\frac{\Delta c}{\theta}$. The importer will choose to visit exporter 1 if and only if $A_1 > A_2$. For an importer with $A_1 > A_2$, the minimum demand parameters for which the importer will be willing to visit exporter 1 will be the values for which

$$S(A_1, c/\theta) + S(A_2, \Delta c/\theta) = F.$$

4.6 Welfare and Trade Costs

This section characterizes the consequences of MPTs and strategic pricing for the welfare gains from trade. Recall, that the importer's demand function from exporter i is $q(A_i, p_C) = \left(\frac{A\theta}{p_C}\right)^{\frac{1}{1-\theta}}$ with surplus $S(A_i, p_C) = \kappa A_i^{\frac{1}{1-\theta}} p_C^{-\frac{\theta}{1-\theta}}$, where p_C is the delivered price to the importer and $\kappa = \theta^{\frac{\theta}{1-\theta}}(1-\theta) > 0$. As in section 4, we maintain that exporter i chooses the price of good j, p_{ij} to maximize $\pi_(A_j, p_{ij}) = (p_j - c_{ij})q(A_j, p_C)$, $i, j \in \{1, 2\}$, where trade costs are assumed to be paid by the importer. The profit maximization problem is subject to the constraint that the importer does not deviate by incurring the fixed trade cost F and purchase good $j \in \{1, 2\}$ from seller $k \neq i$, $S(A_j, p_{Cij}) \geq S(A_j, p_{Ckj}) - F$.

We simplify with the following symmetry assumptions:

Assumption 1 *Symmetry of Costs and Demands*
(i) The unit cost of seller i for good j satisfy $c_{ii} = c$ and $c_{ij} = \Delta c$ for $i = \{1, 2\}$ and $i \neq j$.
(ii) The distribution function for demand shocks $G(A_1, A_2)$ is symmetric, $G(x, y) = G(y, x)$ for $x, y \in [A_{min}, A_{max}]$.

The symmetry of the distribution of shocks allows us to focus on realizations were $A_1 \geq A_2$ and focus on the transactions of seller 1. We examine the cases of per unit and iceberg trade costs in turn.

4.6.1 Importer Surplus and Profits with Trade Costs

Lemma 1 *Per Unit Trade Costs*
Assume trade costs of t per unit paid by the importer.

(i) In the absence of a binding no-deviation constraint, the exporter's price of good j will be $p_j = \frac{c_{1j}+t(1-\theta)}{\theta}$ and the consumer price $p_{Cj} = \frac{c_{1j}+t}{\theta}$ for $j \in \{1,2\}$. The effect of changes in trade costs on importer surplus and exporter profits for given demand parameter A will be

$$\frac{dS(A_j,p_{Cj})}{dt} = -\frac{q(A_j,p_{Cj})}{\theta} \qquad \frac{dS(A_j,p_{Cj})}{dF} = -1 \qquad \frac{d\pi(A_j,t)}{dt} = -q(A_j,p_{Cj}) \quad j \in \{1,2\} \tag{12}$$

(ii) If the no deviation constraint binds, the exporter's price of good 2 satisfies $S(A_2,p+t) = S\left(A_2,\frac{c+t}{\theta}\right) - F$. The effect of trade costs on importer surplus and exporter profits

$$\frac{dS(A_2,p+t)}{dt} = -\frac{q(A_2,\frac{c+t}{\theta})}{\theta} \qquad \frac{dS(A_2,\frac{c+t}{\theta})}{dF} = -1 \tag{13}$$

$$\frac{d\pi(A_2,t)}{dt} = \frac{\lambda q(A_2,\frac{c+t}{\theta})}{\theta} - q(A_2,p_C) \qquad \frac{d\pi(A_2,t)}{dF} = \lambda$$

where $\lambda = 1 + \frac{(p-\Delta c)\frac{\partial q(A_2,p_C)}{\partial p_C}}{q(A_2,p_C)} \in [0,1]$.

In the case where the seller's price is unconstrained, the effect of a change in unit costs on importer welfare is proportional to the quantity of imports and the effect of fixed costs comes directly out of importer surplus. In this case the welfare effect can be measured by summing the quantity of imports for the effect of changes in t and by summing the number of visits to seller 1 for changes in F.

In the case where the seller's price is constrained by the no deviation constraint, the level of welfare must satisfy $S(A_2,p) = S(A_2,\frac{c+t}{\theta}) - F$. The effect of a change in unit trade costs is then $\frac{\partial S(A_2,\frac{c+t}{\theta})}{\partial t}$, which is proportional to the quantity that the importer would have bought if it deviated to the lower cost seller. Since $\frac{c+t}{\theta} < p_C$ when the constraint binds and $\theta < 1$, using the observed quantity of imports for the non-core good would understate the importer's loss in the case of an increase in t (or understate the importer's gain in the case of a decrease in t). This difference is due to the strategic effect: increases (decreases) in t and F benefit the exporter (importer) by relaxing (tightening) the no deviation constraint on the exporter's price. For the exporter, an increase in t has conflicting effects because it reduces the demand for its product but relaxes the no deviation constraint. Note that $\lambda \in [0,1]$ reflects the fact that the exporter's price is constrained below the monopoly level, with $\lambda = 0$ when $p = \frac{\Delta c+t(1-\theta)}{\theta}$ and $\lambda = 1$ at $p = \Delta c$. Similarly, an increase in F benefits the exporter by relaxing the no deviation constraint and a decrease in F harms the exporter, with the magnitude depending on the extent to which the constraint binds.

Lemma 2 *Iceberg Trade Costs*

With trade costs τp paid by the importer, where p is the producer price,
(i) When the no deviation constraint does not bind, the exporter chooses p to maximize $(p-c)q(A,\tau p)$. The price that maximizes producer profit in this case is $\frac{c_{1j}}{\theta}$ and the consumer price $\frac{\tau c_{1j}}{\theta}$. The impact of a change in trade costs on welfare and profits is

$$\frac{dS(A_j,p_{Cj})}{d\tau} = -p_j q(A_j p_{Cj}) \qquad \frac{dS(A_j,p_{Cj})}{dF} = -1 \qquad \frac{d\pi(A_j,t)}{d\tau} = -\frac{p_j q(A_j,p_{Cj})}{\tau} \tag{14}$$

(ii) If the no deviation constraint binds, the exporter's price satisfies $S(A, p+t) = S\left(A, \frac{\tau c}{\theta}\right) - F$. *The effect of trade costs on importer surplus and exporter profits is*

$$\frac{dS(A_2, \tau p_{C2})}{d\tau} = -p_2 q(A_2, \frac{\tau c}{\theta}) \qquad \frac{dS(A_2, \tau p_{C2})}{dF} = -1 \tag{15}$$
$$\frac{d\pi(A_2, \tau)}{d\tau} = \mu p_2 q(A, \frac{\tau c}{\theta}) - q(A_2, \tau p_2) \qquad \frac{d\pi(A_2, \tau)}{dF} = \mu$$

where $\mu = 1 + \frac{(p-\Delta c)\frac{\partial q(A_2, \tau p)}{\partial p}}{q(A_2, \tau p)} \in [0, 1]$.

$S(A_2, \tau p) = S(A_2, \frac{\tau c}{\theta}) - F$ To see how the exporter's factory-gate price p must change when the iceberg cost τ changes, we take the total derivative of both sides with respect to τ. In this model, the derivative of surplus with respect to the delivered price is just the negative of the quantity demanded. Applying the chain rule to both sides gives us the following equation. $-q(A_2, \tau p) \cdot (p + \tau \frac{dp}{d\tau}) = -q(A_2, \frac{\tau c}{\theta}) \cdot (\frac{c}{\theta})$ Next, we want to isolate the derivative of the exporter's price, $dp/d\tau$. We divide both sides by the negative quantity $-q(A_2, \tau p)$. $p + \tau \frac{dp}{d\tau} = \frac{c}{\theta} \frac{q(A_2, \frac{\tau c}{\theta})}{q(A_2, \tau p)}$ Now we subtract p from both sides, and then divide the entire equation by τ to get the derivative completely by itself. This division is exactly where the $1/\tau$ factor is born. $\frac{dp}{d\tau} = \frac{c}{\theta\tau} \frac{q(A_2, \frac{\tau c}{\theta})}{q(A_2, \tau p)} - \frac{p}{\tau}$ Now that we know how the exporter's price responds to shipping costs, we can look at the profit equation. The exporter's profit is their markup multiplied by the quantity sold. $\pi = (p - \Delta c) \cdot q(A_2, \tau p)$ We take the total derivative of this profit equation with respect to τ using the product rule. The change in profit equals the change in the price multiplied by the quantity, plus the markup multiplied by the change in the quantity. $\frac{d\pi}{d\tau} = \frac{dp}{d\tau} \cdot q(A_2, \tau p) + (p - \Delta c) \cdot \frac{d}{d\tau}[q(A_2, \tau p)]$ Notice that the change in the delivered price inside that last quantity term is exactly $p + \tau(dp/d\tau)$. We already solved for that specific cluster of terms a few steps ago, which was equal to $(c/\theta) \cdot q(A_2, \tau c/\theta)/q(A_2, \tau p)$. If we let q_P represent the derivative of quantity with respect to the delivered price, we can substitute our earlier finding directly into the second half of the profit equation. We can also substitute our long $dp/d\tau$ formula into the first half. $\frac{d\pi}{d\tau} = \left[\frac{c}{\theta\tau} \frac{q(A_2, \frac{\tau c}{\theta})}{q(A_2, \tau p)} - \frac{p}{\tau}\right] q(A_2, \tau p) + (p - \Delta c) \cdot q_P \cdot \left[\frac{c}{\theta} \frac{q(A_2, \frac{\tau c}{\theta})}{q(A_2, \tau p)}\right]$ When we multiply the $q(A_2, \tau p)$ through the first bracket, the denominator cancels out on the first term, leaving us with a clean expression for the first half of the equation. $\frac{d\pi}{d\tau} = \frac{c}{\theta\tau} q(A_2, \frac{\tau c}{\theta}) - \frac{p}{\tau} q(A_2, \tau p) + (p - \Delta c) \cdot q_P \cdot \frac{c}{\theta} \frac{q(A_2, \frac{\tau c}{\theta})}{q(A_2, \tau p)}$ Our final step is to group the terms together and introduce the paper's μ variable. We can pull the competitor's term out of both the first and third pieces of that equation. Notice that to factor out $(c/\theta\tau)$, we have to multiply the third piece by τ to keep the algebra balanced. $\frac{d\pi}{d\tau} = \frac{c}{\theta\tau} q(A_2, \frac{\tau c}{\theta}) \left[1 + (p - \Delta c) \frac{q_P \cdot \tau}{q(A_2, \tau p)}\right] - \frac{p}{\tau} q(A_2, \tau p)$ The paper defines μ based on the derivative of quantity with respect to the factory price p. By the chain rule, the derivative with respect to p is exactly $q_P \cdot \tau$. Therefore, the entire bracketed section is perfectly identical to the paper's definition of μ. Substituting μ in place of the brackets gives us our final, corrected formula. $\frac{d\pi}{d\tau} = \mu \frac{c}{\theta\tau} q(A_2, \frac{\tau c}{\theta}) - \frac{p}{\tau} q(A_2, \tau p)$

With iceberg trade costs, the impact of the change in variable costs on consumer welfare is measured by the value of trade at fob prices when the no deviation constraint is not binding. When the no deviation constraint binds, the appropriate measure is the fob value of imports if the importer deviates to the more efficient seller. As in the case with per unit trade costs, the observed fob value understates

the welfare change due to the strategic effect when the no deviation constraint binds.

4.6.2 Aggregate Welfare Effects with Per Unit Trade Costs

Importers visiting seller 1 consist of three groups: group L that purchase both goods from seller 1 at the monopoly price, group I purchases both goods from seller 1 but receives a discount on good 2, and group H that only purchases good 1 from seller 1. For the case of per unit trade costs, we can sum over the welfare effects for importers in each group using the individual importer effects from (12) and (13). Since importers at the margin are indifferent between groups, the effects of importers switching groups as a result of trade cost changes will be of second order. The case with iceberg trade costs can be similarly analyzed using (14) and (15).

The following result characterizes the sets Ω_i consisting of the values of (A_1, A_2) contained in group $i \in \{L, I, H\}$.

Lemma 3 *Define $\underline{A}_2$ to be the value satisfying $S\left(A_2, \frac{\Delta c}{\theta}\right) = S\left(A_2, \frac{c+t}{\theta}\right) - F$ and $\bar{A}_2$ to be the value satisfying $S\left(A_2, \Delta c + t\right) = S\left(A_2, \frac{c+t}{\theta}\right) - F$.*

(i) $\Omega_L = \{(A_1, A_2) | A_2 \leq \underline{A}_2\}$ and $S\left(A_1, \frac{c+t}{\theta}\right) + S\left(A_2, \frac{\Delta c+t}{\theta}\right) - F \geq 0$.

(ii) $\Omega_I = \{(A_1, A_2) | A_2 \in \left[\underline{A}_2, \bar{A}_2\right]\}$

(iii) $\Omega_H = \{(A_1, A_2) | A_2 \geq \bar{A}_2\}$

Proof. (**Lemma 3**) The set Ω_L consists of importers whose return from buying from seller 1 is non-negative and who prefer buying good 2 from seller 1 at the monopoly price $\frac{\Delta c}{\theta}$ to incurring the fixed cost of visiting seller 2 and paying its monopoly price. Since $S\left(A_2, \frac{\Delta c+t}{\theta}\right) - S\left(A_2, \frac{c+t}{\theta}\right)$ is decreasing in A_2, the latter condition will hold for $A_2 \leq \underline{A}_2 = \left(\frac{F}{\gamma}\left(\left(\frac{c+t}{\theta}\right)^{\frac{\theta}{\theta-1}} - \left(\frac{\Delta c+t}{\theta}\right)^{\frac{\theta}{\theta-1}}\right)\right)^{\theta-1}$.

The set Ω_I consists of importers whose return from buying from seller 1 is non-negative and for whom there is a price p_2 such that $S\left(A_2, \frac{\Delta c+t}{\theta}\right) - F = S\left(A_2, p_2 + t\right)$, which will be satisfied for all $A_2 \in \left[\underline{A}_2, \bar{A}_2\right]$. Similarly, the set Ω_H contain importers that engage in single product transactions for the core product of each seller. For $A \geq \bar{A}_2$ the importer will prefer purchasing good 2 from seller 2 to buying from seller 1 at seller 1's marginal cost. Thus,

$$\underline{A}_2 = \left(\frac{F}{\gamma}\left(\left(\frac{c+t}{\theta}\right)^{\frac{\theta}{\theta-1}} - \left(\frac{\Delta c+t}{\theta}\right)^{\frac{\theta}{\theta-1}}\right)\right)^{\theta-1} \quad \bar{A}_2 = \left(\frac{F}{\gamma}\left(\left(\frac{c+t}{\theta}\right)^{\frac{\theta}{\theta-1}} - (\Delta c+t)^{\frac{\theta}{\theta-1}}\right)\right)^{\theta-1} \tag{16}$$

■

Discussion. For an importer in Ω_L, the effect of changes in transport costs is

$$dW(A_1, A_2) = -\frac{1}{\theta}\left[q\left(A_1, \frac{c+t}{\theta}\right) + q\left(A_2, \frac{\Delta c+t}{\theta}\right)\right] dt - dF \quad \text{for } (A_1, A_2) \in \Omega_L \tag{17}$$

Summing over all the low demand importers, the change in welfare from a change in t will be proportional to the quantity of imports and the effect of a change in F will be proportional to the number of importers. In this sense, among low demand firms gains from trade liberalization are analogous to

those in conventional, workhorse models in which the gains are proportional to trade flows up to a trade cost elasticity parameter (Arkolakis, Costinot, and Rodríguez-Clare, 2012).

Since importers in group I are indifferent between buying from exporter 1 or deviating to exporter 2, the effect of trade costs on welfare for groups I and H is

$$dW(A_1, A_2) = -\frac{1}{\theta}\left[q\left(A_1, \frac{c+t}{\theta}\right) + q\left(A_2, \frac{c+t}{\theta}\right)\right] dt - 2dF \quad \text{for } (A_1, A_2) \in \Omega_I \cup \Omega_H \tag{18}$$

As discussed in relation to Lemma 1, measuring the gains from transport cost reductions using observed import quantities and visits to seller to measure the gains from trade cost changes will understate the effect of changes for importers in group I due to the strategic effect.

The total welfare effects are obtained by integrating (17) and (18) using the distribution of importer demands for the respective groups. The change in the strategic effect is obtained by integrating over importers in group I,

$$\int_{\Omega_I} \int \left[-\frac{1}{\theta}\left(q\left(A_2, \frac{c+t}{\theta}\right) - q(A_2, p(A_2)+t)\right) dt - dF\right] g(A_1, A_2) dA_1 dA_2 \tag{19}$$

Importers in group I generate additional gains from trade liberalization because of strategic pricing. A reduction in trade costs makes switching to SPTs more attractive. Accordingly, firms with an active outside option (even though it is unrealized) capture additional discounts from exporter 1. This additional welfare effect is similar in spirit to the "pro-competitive" gains from trade (Krugman, 1979; Markusen, 1981; Atkeson and Burstein, 2008; Edmond, Midrigan, and Xu, 2015) but arise from the ability of importers to purchases bundles of goods from a single supplier or sourcing input from independent suppliers.

5 Testable Implications

Implication 1A. Two importers have the same demand for a supplier's core product, and both buy it from that supplier ($q_i \geq q_i^0$). The one with lower demand for a peripheral product is more likely to add it to the order. The other buys it separately from the supplier for which it is the core product.

We compare the quantity of a peripheral product in an MPT with its quantity in an SPT. We use SPTs from another importer that bought a similar quantity of the core product or from the same importer at another date.

Implication 1B. Take two importers of the same size who buy the same products and spend the same amount in total. One buys them together from one supplier. The other buys each product separately from its core supplier. We expect the first importer to spend a larger share on its supplier's core product and a smaller share on each other product.

We look for SPTs of the same products that add up to about the same value as the MPT. We use purchases by another importer of similar size or by the same importer at another date.

Implication 2A. Two importers buy the same supplier's core product. The one that adds a peripheral product to the order pays more per unit for it than the one that buys it separately from its core supplier.

We compare the unit values of peripheral products in MPTs with those of the same products sold separately by specialists. We use past unit values to choose which specialists to include.

Implication 2B. Among importers with the same demand for the core product, we expect those who bundle a peripheral product to spend less on it than those who buy it separately from its core supplier. The smaller quantity would have to outweigh the higher unit price.

We repeat the comparison in Implication 1A using spending on the peripheral product.

Implication 3. Starting from equal MPT shares, a small rise in one supplier's relative trade costs leaves it with the lower MPT share.

We estimate how often transactions are multiproduct at different supplier trade costs.

5.1 Matching MPTs to SPTs

Specialists. We use past sales to identify an exporter's specialties. Let $V_{ep\tau}$ be the value of exporter e's observed sales of HS6 product p in year τ, and $V_{e\tau}$ be its total observed sales value that year. The share of its sales from product p over the previous three years is

$$\text{Share}_{ept}^{\text{trail}} = \frac{\sum_{\tau=t-3}^{t-1} V_{ep\tau}}{\sum_{\tau=t-3}^{t-1} V_{e\tau}}. \tag{20}$$

An exporter is a *portfolio specialist* in a product if, over those three years, the product ranked among its top three by sales value, accounted for at least 15 percent of its sales, and appeared in at least ten of its SPTs.

For a given purchase, a *qualifying specialist* is a specialist that meets these criteria even after we remove the buyer's past purchases, or that sold the product in SPTs to at least three other importers within 30 days of the purchase.

Product roles. The core product is the specialty in the MPT with the largest share of the exporter's sales over the previous three years. Any other specialty in the MPT is a secondary specialty. The remaining products are peripheral. The core can be a small line in the MPT because we assign roles using past sales.[16]

In 64.0 percent of MPTs with a specialty core, the importer accounted for at least half of the exporter's sales over the previous three years. We assign roles again using only the exporter's sales to other buyers. We keep an MPT only if every product keeps its role.

[16] Choosing the largest line as the core would make the other lines smaller by construction and build part of the pattern in Implications 1A and 1B into the definition.

Sample. Of the 10,473,809 MPTs in 2014–2020, 3,179,927 have a specialty core. Of those, 1,278,646 pass the check on product roles. In our sample, 81 percent of transactions are definitive imports and 95 percent report a mode of transport. We keep definitive imports with a recorded mode of transport. This leaves 972,397 MPTs. We impose further restrictions when valuing products at reference prices.

Synthetic MPTs. For an MPT i by importer m from exporter e, let c be the core product and $\mathcal{P}_i$ the set of required products. We drop the MPT if m bought a required product from a specialist in that product in any transaction within 30 days of the MPT.

We look for a different importer j whose annual import value in the year of the MPT is within a factor of $\exp(0.5) \approx 1.65$ of m's. We apply the following rules to j's purchases within 30 days of the MPT.

1. Importer j bought the core product c in an SPT from a specialist in c, possibly exporter e.
2. Importer j bought each required product $p \in \mathcal{P}_i$ in an SPT.
3. Importer j bought no required product in an MPT or in an SPT from exporter e.
4. The SPTs used for the match are definitive imports shipped by the same mode as the MPT.

For Implications 1A and 2B, we also match on core quantity. For Implication 1B, we match on total spending on the products in the MPT. For each required product, we take the SPT nearest in date to the MPT that records quantity in the same unit. We choose the importer whose purchases come closest to the MPT in date and in the quantity or spending used for the test. We break ties using the lowest importer ID. That importer's SPTs of the core and required products form a *synthetic MPT*. We can use the same importer for more than one match.

6 Concluding Remarks

This paper establishes that a large fraction of trade is conducted through multiproduct transactions: transactions between an exporter and importer such that the two parties are contracting over multiple goods within a single transaction. Though common, MPTs typically comprise a small number of goods. In aggreagte, MPTs account for 31 percent of the value of Mexican imports in our data.

We document that sourcing patterns vary significantly across firm size: small firms typically purchase multiple goods from a single producer and bundle goods within individual transactions. Conditional on firm size, larger purchases of a product are more likely to be made in single-product transactions. On average, large firms are more likely to source individual goods from distinct suppliers, while small firms are more likely to source multiple goods from a single supplier.

These patterns have important implications for the gains from trade. A reduction in trade costs not only enables firms to source individual inputs from their lowest-cost suppliers, expanding supplier networks and lowering downstream production costs, but also changes the organization of international transactions. As firms substitute away from multiproduct transactions toward single-product transactions, suppliers face stronger competitive pressure for individual products and reduced scope to exercise

market power over peripheral goods. Trade liberalization therefore generates welfare gains through two complementary channels: improved allocation of products across suppliers and a pro-competitive reduction in the prices charged within buyer–seller relationships. The margin between multiproduct and single-product transactions is thus central to understanding how trade costs shape both firm sourcing decisions and the welfare gains from international integration.

References

Amador, J., and L. D. Opromolla (2013): "Product and destination mix in export markets," *Review of World Economics (Weltwirtschaftliches Archiv)*, 149(1), 23–53.

Antràs, P., T. C. Fort, and F. Tintelnot (2017): "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 107(9), 2514–64.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012): "New Trade Models, Same Old Gains?," *American Economic Review*, 102(1), 94–130.

Arkolakis, C., S. Ganapati, and M.-A. Muendler (2021): "The Extensive Margin of Exporting Products: A Firm-Level Analysis," *American Economic Journal: Macroeconomics*, 13(4), 182–245.

Arkolakis, Costas, F. H., and Y. Miyauchi (2025): "Production Network Formation, Trade, and Welfare," NBER Working Paper 30954.

Atkeson, A., and A. Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98(5), 1998–2031.

Baldwin, J., and W. Gu (2009): "The Impact of Trade on Plant Scale, Production-Run Length, and Diversification," in *Producer Dynamics: New Evidence from Micro Data*. University of Chicago Press.

Bastos, P., and J. Silva (2012): "Networks, firms, and trade," *Journal of International Economics*, 87(2), 352–364.

Bernard, A., J. Jensen, S. Redding, and P. Schott (2012): "The Empirics of Firm Heterogeneity and International Trade," *Annual Review of Economics*, 4(1), 283–313.

Bernard, A., A. Moxnes, and K. H. Ulltveit-Moe (2018): "Two-Sided Heterogeneity and Trade," *The Review of Economics and Statistics*, 100(3), 424–439.

Bernard, A., S. Redding, and P. Schott (2011): "Multiproduct Firms and Trade Liberalization," *The Quarterly Journal of Economics*, 126(3), 1271–1318.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2018): "Global Firms," *Journal of Economic Literature*, 56(2), 565–619.

Bernard, A. B., S. J. Redding, and P. K. Schott (2010): "Multiple-Product Firms and Product Switching," *American Economic Review*, 100(1), 70–97.

Berthou, A., and L. Fontagné (2013): "How do Multiproduct Exporters React to a Change in Trade Costs?," *The Scandinavian Journal of Economics*, 115(2), 326–353.

Chaney, T. (2014): "The Network Structure of International Trade," *American Economic Review*, 104(11), 3600–3634.

Demir, B., A. C. Fieler, D. Y. Xu, and K. K. Yang (2024): "O-Ring Production Networks," *Journal of Political Economy*, 132(1), 200–247.

Dhingra, S. (2013): "Trading Away Wide Brands for Cheap Brands," *American Economic Review*, 103(6), 2554–84.

Eaton, J., D. Jinkins, J. Tybout, and D. Y. Xu (2022): "Two-sided Search in International Markets," Working papers, U.S. Census Bureau, Center for Economic Studies.

Eckel, C., and J. P. Neary (2010): "Multi-Product Firms and Flexible Manufacturing in the Global Economy," *The Review of Economic Studies*, 77(1), 188–217.

Edmond, C., V. Midrigan, and D. Y. Xu (2015): "Competition, Markups, and the Gains from International Trade," *American Economic Review*, 105(10), 3183–3221.

Feenstra, R., and H. Ma (2008): "Optimal choice of product scope for multiproduct firms," in *The Organization of Firms in a Global Economy*, ed. by E. Helpman, D. Marin, and T. Verdier. Harvard University Press.

Huneeus, F. (2018): "Production Network Dynamics and the Propagation of Shocks," Working Paper.

Iacovone, L., and B. S. Javorcik (2010): "Multi-Product Exporters: Product Churning, Uncertainty and Export Discoveries," *The Economic Journal*, 120(544), 481–499.

Krugman, P. R. (1979): "Increasing returns, monopolistic competition, and international trade," *Journal of International Economics*, 9(4), 469–479.

Lopresti, J. (2016): "Multiproduct firms and product scope adjustment in trade," *Journal of International Economics*, 100(C), 160–173.

Manova, K., and Z. Yu (2017): "Multi-product firms and product quality," *Journal of International Economics*, 109(C), 116–137.

Markusen, J. R. (1981): "Trade and the Gains from Trade with Imperfect Competition," *Journal of International Economics*, 11(4), 531–551.

Mayer, T., M. Melitz, and G. Ottaviano (2014): "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 104(2), 495–536.

Qiu, L. D., and W. Zhou (2013): "Multiproduct firms and scope adjustment in globalization," *Journal of International Economics*, 91(1), 142–153.

Rauch, J. E. (1999): "Networks versus markets in international trade," *Journal of International Economics*, 48(1), 7–35.

Thomas, C. (2011): "Too Many Products: Decentralized Decision Making in Multinational Firms," *American Economic Journal: Microeconomics*, 3(1), 280–306.

Timoshenko, O. (2015): "Product switching in a model of learning," *Journal of International Economics*, 95(2), 233–249.