

When Tariffs Raise Markups: Import Competition and Local Market Power*

David Argente

Yale University and NBER

Esteban Méndez

Central Bank of Costa Rica

Diana Van Patten

Yale University and NBER

June 29th, 2026

Abstract

Trade policy can affect consumers not only through border prices but also by shifting the market power of domestic retailers in segmented local markets. We study this channel using Costa Rica’s elimination of its *de minimis* exemption, which raised the delivered price of direct-to-consumer imports while leaving domestic retailers’ cost structure unchanged. We link shipment-level customs records for individual parcels to the universe of retail electronic invoices and to customers’ locations and earnings. These data allow us to construct *retailer-specific catchment areas* and observe domestic and foreign expenditure shares in the relevant markets—the collection of neighborhoods each retailer actually serves. We find that more exposed retailer-product cells raise prices and sales, while import and purchase prices, wages, and employment do not respond, consistent with higher markups rather than cost pass-through. We then develop a catchment-area model with variable markups and strategic complementarities, in which market shares are defined over retailer-specific catchment areas, and discipline its key elasticities using *observed* share changes instrumented by the tariff shock. Guided by the model, we show that price responses are sharply nonlinear in market shares, steepening in the upper tail. Crucially, *de minimis* exposure is highest in richer neighborhoods (where direct importing was common), but retail concentration is systematically higher in poorer neighborhoods where households have the fewest local alternatives. Therefore, a reform that appears progressive when judged by exposure can be regressive in equilibrium. Welfare losses are elusive on average but large in low-competition, lower-income areas. The results show how the *de minimis* reform—and tariffs more broadly—can amplify local market power and shift incidence toward consumers in low-competition local markets.

Keywords: *de minimis*, market power, tariffs, variable markups, spatial competition.

*Corresponding author: diana.vanpatten@yale.edu. We thank Fernando Alvarez, David Atkin, Johannes Boehm, Mathilde Bombardini, Ariel Burstein, Lorenzo Caliendo, Ben Faber, Amit Khadelwal, Simon Mongey, Ezra Oberfield, Andrés Rodríguez-Clare, Esteban Rossi-Hansberg, Todd Schoellman, Pete Schott, Mike Waugh, and participants at various seminars and conferences for helpful comments. The views expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Costa Rica (BCCR). Results have been reviewed by the BCCR to ensure no confidential information is disclosed.

1 Introduction

Small-parcel imports purchased directly by consumers have expanded rapidly, propelled by global marketplaces, cheaper logistics, and the diffusion of digital payments. In many countries, these shipments receive preferential treatment under *de minimis* regimes—thresholds below which tariffs are waived. While *de minimis* policies are often framed as administrative conveniences for low-value items, they can be quantitatively important in aggregate (Fajgelbaum and Khandelwal, 2024) and can provide meaningful competitive pressure on domestic retailers by strengthening households’ foreign outside option. This exception has also become the focus of active policy debate: in 2025, the White House announced the suspension of the *de minimis* provision, which was subsequently continued in 2026, prompting renewed discussion in the United States and abroad about eliminating such exemptions.¹

These debates raise a broader question about the incidence of online retail and of selective tariff exemptions. Direct-to-consumer online purchases can expand households’ choice sets and may be especially valuable in markets where local retail options are limited. At the same time, the consumer benefit of a tariff exemption need not be determined only by who uses the exempted channel directly. Under imperfect competition, a selective exemption can also discipline domestic firms by changing the outside option available to consumers; conversely, removing the exemption can relax competitive pressure and increase domestic markups.

This paper exploits the elimination of Costa Rica’s *de minimis* exemption as an import-competition shock that raises the price of direct-to-consumer imports and thereby increases domestic retailers’ effective market power in exposed local markets, while leaving domestic retailers’ marginal costs unchanged. In May 2021, Costa Rica adopted the Central American Uniform Customs Code to harmonize customs procedures within the region. As an unintended consequence, the country eliminated its tariff exemption for non-commercial shipments valued at or below \$500, increasing the average tariff-and-tax burden on these shipments by 8 percentage points (about 57 percent) and reducing the value of consumer imports by roughly 50 percent. The reform therefore provides a setting to study the incidence of removing a selective tariff exemption and, more broadly, how making households’ online outside option more costly reshapes competition and consumer incidence in imperfectly competitive local retail markets.

¹See The White House (2025, 2026).


We study this shock in a data-rich setting that combines individual-level information on de minimis imports by product with detailed retail transaction-level data on the prices, sales, and varieties sold by the universe of domestic retailers. Crucially, the data also link consumers to their residence neighborhoods and to the retailers where they shop. We use these shopping links to recover the geography of each retailer’s customer base. For each retailer, we construct a sales-weighted distribution over the neighborhoods from which its customers originate; this distribution defines the retailer’s catchment area. This allows us to measure product-specific domestic and foreign shares in the markets a retailer actually serves, rather than imposing administrative boundaries, fixed-distance radii, or aggregate national shares. We validate this construction using two alternative approaches—a gravity-based reconstruction of retailer–neighborhood links and a proxy based on employees’ residence neighborhoods—and show that both the main reduced-form results and the estimated markup-elasticity schedule are robust across these market definitions. The catchment-area structure is also important for incidence because retailers often price over the set of customers they serve rather than over a single neighborhood. As a result, direct exposure in one neighborhood can affect the common price paid by other neighborhoods in the same retailer catchment, while less-exposed neighborhoods can dampen the price response. This pricing-market pooling makes direct exposure alone an incomplete proxy for incidence.

Building on these catchment-area measures, we quantify adjustments in markups and in the breadth of retailers’ assortments by leveraging variation in tariff changes across products and spatial variation in pre-reform de minimis penetration across the neighborhoods served by each retailer. Specifically, a retailer’s *ex ante* exposure depends on (i) the category-specific increase in the tariff burden on de minimis imports and (ii) the pre-policy reliance on de minimis imports among households in the retailer’s catchment area. This exposure measure therefore combines product-level variation in the policy shock with customer-based variation in the foreign competitive pressure faced by each retailer. Empirically, we find that more exposed retailers increase prices and nominal sales in affected categories, without corresponding increases in input costs, wages, or employment. At the same time, domestic assortment expands while foreign variety contracts.

Importantly, the largest price increases do not occur in the *ex ante* most exposed locations—those where residents relied most on de minimis importing prior to the reform. Instead, price responses are strongest in exposed locations where high baseline

product-specific retail concentration allows the reduction in import competition to translate into substantial retailer market power, enabling retailers to raise prices sharply as the foreign outside option becomes more expensive. This distinction is key because the direct users of the exemption are not necessarily the consumers who ultimately bear the largest equilibrium price effects. The diagram in Figure 1 summarizes these dynamics. Notably, neighborhoods with high *ex ante* exposure to the de minimis regime tend to be higher-income, whereas neighborhoods with high product-specific retail concentration tend to be lower-income. As a result, a policy that appears progressive when judged solely by exposure can be regressive in equilibrium: the largest price increases arise in highly concentrated markets, which are more prevalent in lower-income areas. This pattern reflects pronounced nonlinearity in price responses at the upper tail of the market-share distribution, where markups are most sensitive to changes in competitive pressure.

Figure 1: Exposure and Local Concentration



Notes: The figure summarizes two dimensions of heterogeneity driving our results. On the left, neighborhoods differ in pre-policy de minimis exposure per product: high in neighborhood 1, moderate in 2, and zero in 3. On the right, neighborhoods differ in local retail competition for a good. The policy's impact depends on these two margins.

These facts motivate a model of spatially segmented local markets with variable markups à la [Atkeson and Burstein \(2008\)](#) with strategic complementarities, in which market shares are defined over retailer-product catchment areas and markup elasticities increase with local market share. Unlike constant-markup, ACR-style sufficient-

statistic frameworks, this structure implies that a policy shock that raises the delivered price of the foreign outside option—without commensurately shifting domestic marginal costs—generates first-order markup responses and sharply heterogeneous (and thus distributionally important) local price effects. In this sense, competitive effects that can appear *elusive* in aggregate gravity environments (Arkolakis, Costinot, Donaldson, and Rodríguez-Clare, 2019) become measurable and economically important once local concentration and nonlinear markup elasticities are incorporated.

Disciplined by rich data on prices, quantities, and market structure, and estimating the key demand elasticities with an IV strategy, the model closely replicates the sharp empirical nonlinearity in markup elasticities across the market-share distribution. Thus, the paper also provides a direct validation of the Atkeson-Burstein framework in an unusually data-rich environment where the main outcomes of interest are observed. The implied elasticities are in line with Broda and Weinstein (2010). Quantitatively, we find that, while the median welfare effect is small, there is substantial heterogeneity: welfare losses are concentrated in low-competition, lower-income neighborhoods, even though *de minimis* importing is more common among higher-income households with access to more retail options.

Related Literature and Contribution. The paper contributes to several strands of the international trade literature that study how trade costs and trade policy shape consumer outcomes and market power in the domestic economy. A first related literature studies the rise of direct-to-consumer trade. As low-value parcels are often imperfectly captured in standard trade statistics, empirical work on direct-to-consumer imports has only recently become feasible. Fajgelbaum and Khandelwal (2024) use novel carrier and customs data to quantify the aggregate and distributional importance of *de minimis* imports in the United States, emphasizing the role of the threshold and the incidence of tariff and processing costs. Argente, Mendez, and Van Patten (2025) leverage consumer imports data to show how social ties spread demand for new products among consumers. The paper shows how retailers learn from early consumers who import and expand their offerings, amplifying access to global varieties through local demand externalities. Finally, Du, Fang, Ma, Tang, and Yu (2025) document that *de minimis* e-commerce rose during the U.S.-China trade war, alongside the tariff increases. Relative to this emerging literature, we study the elimination of a *de minimis* exemption as an increase in the delivered price of direct-to-consumer imports that leaves retailers’ cost structure unchanged, to trace the responses in domestic

retail prices, markups, and assortments. We emphasize a quantitatively meaningful nonlinearity: price responses steepen at high market shares, so the largest increases occur where retailers have the greatest market power in the relevant product market. This also connects the *de minimis* debate to a broader question about online retail: whether direct access to remote sellers disciplines local retail market power, and how the answer depends on the pricing behavior and geographic footprint of domestic retailers.

Second, we speak to the classic question of whether trade liberalization has pro-competitive effects that reduce markups and distortions in domestic markets. A large theoretical literature allows markups to vary with competitive pressure, including demand-based models with variable markups and strategic complementarities (e.g., [Krugman, 1979](#); [Melitz and Ottaviano, 2008](#); [Atkeson and Burstein, 2008](#)). In influential work, [Arkolakis, Costinot, Donaldson, and Rodríguez-Clare \(2019\)](#) show that within a broad class of gravity models, pro-competitive effects are often *elusive* in aggregate welfare terms: moving from constant- to variable-markup environments does not amplify gains from trade and may even dampen them once the behavior of foreign markups is accounted for. Our paper complements this insight by shifting attention from sector-level competition in aggregate gravity environments to *spatially segmented local markets*. With nonlinear markup elasticities that rise sharply with local market share, an increase in the cost of the foreign outside option generates first-order and sharply heterogeneous markup responses, making pro-competitive effects empirically salient once local concentration and spatial segmentation are incorporated.

Third, our analysis relates to empirical work on how trade shocks affect prices, pass-through, and markups. [Amiti, Itskhoki, and Konings \(2019a\)](#) develop a framework to estimate strategic complementarities in price setting and recover markup elasticities using Belgian manufacturing firm-product data. Together with [Amiti, Itskhoki, and Konings \(2014\)](#), this work shows that market shares are central for markup adjustment. Our paper builds on this insight by bringing the markup-elasticity logic to spatially segmented local markets. In our setting, as the reform leaves domestic retailers' marginal costs unchanged, we can straightforwardly separate markup responses from cost pass-through. Moreover, rather than relying on national or sectoral firm shares, we construct retailer-specific catchment areas and observe both domestic and foreign expenditure shares in the relevant local markets. This retailer-specific market delineation reveals a substantial mass of retailer-product pairs with high market shares, allowing us to estimate the markup-elasticity schedule

in the upper tail, where the mechanism by [Amiti, Itskhoki, and Konings \(2019a\)](#) predicts adjustment should be strongest. This upper-tail variation is key, as it turns the import competition shock into large and uneven local price and welfare effects once the implied markup responses are mapped into local consumer welfare. Contemporaneous work by [Blumenfeld, Patterson, and Vavra \(2026\)](#) also emphasizes that the market-power-relevant notion of shares is customer based: using credit-card data, they focus on the share of each customer’s spending captured by a merchant and show how these measures differ from national shares. Our setting differs in that the de minimis reform provides quasi-experimental variation in competitive pressure, allowing us to use these shares to recover markup elasticities and quantify welfare incidence. Moreover, our gravity- and employee-residence-based alternatives show how similar catchment-area market definitions can be implemented when direct customer links are unavailable.

More broadly, using production data and structural restrictions, studies such as [De Loecker, Goldberg, Khandelwal, and Pavcnik \(2016\)](#) document that trade reforms can coincide with rising measured markups even as marginal costs fall, underscoring the importance of separating cost changes from market-power adjustments. In the context of the recent U.S. trade war, a growing literature finds substantial tariff pass-through to border prices and meaningful effects on trade flows and consumer prices (e.g., [Amiti, Redding, and Weinstein, 2019b](#); [Fajgelbaum, Goldberg, Kennedy, and Khandelwal, 2020](#); [Cavallo, Gopinath, Neiman, and Tang, 2021](#)). Relative to these studies, our setting features a particularly clean import-competition shock: the Costa Rican reform raises the delivered price of de minimis imports while leaving domestic retailers’ marginal costs unchanged. This separation allows us to attribute price responses to changes in *effective market power*—rather than to cost pass-through or input-price shocks—and to connect these responses to local market structure.

Fourth, we build on work studying retail globalization and consumer welfare, most directly [Atkin, Faber, and Gonzalez-Navarro \(2018\)](#), who quantify welfare gains from foreign supermarket entry and decompose them into new varieties/amenities and pro-competitive price reductions at incumbent stores. Their approach connects store entry to the broader trade literature on gains from variety (e.g., [Feenstra, 1994](#); [Broda and Weinstein, 2006](#); [Feenstra and Weinstein, 2017](#)) and to work measuring heterogeneous and distributional consequences of globalization (e.g., [Porto, 2006](#); [Topalova, 2010](#); [Goldberg and Pavcnik, 2007](#); [Fajgelbaum and Khandelwal, 2016](#)). Our setting highlights how the strength of pro-competitive forces depends on local market

structure. Our reform contracts the foreign outside option faced by consumers, and our linked administrative data allow us to measure resulting adjustments in domestic retail markups and assortments. After constructing retailer-specific catchment areas, we identify sharp heterogeneity in these adjustments across markets with different baseline retail concentration, isolating when the outside option meaningfully disciplines incumbent behavior. The latter speaks to [Mongey and Waugh \(2025\)](#), who highlight how heterogeneity in demand elasticities and sorting across firms can generate “pricing inequality” and amplify the welfare consequences of shocks.

Finally, our paper relates to work linking trade policy and market power. [Morlacco \(2019\)](#) studies buyer power in input markets and shows that importers’ market power in foreign sourcing can distort input choices and reduce the gains from trade. [Alviarez, Fioretti, Kikkawa, and Morlacco \(2023\)](#) develop a theory of two-sided market power in firm-to-firm trade, showing how bargaining and buyer–supplier network structure shape markups and pass-through. Related contemporary work by [Adão, Fernandes, Hsieh, and Quintero \(2025\)](#) links trade policy and market power through importer concentration. We study market power on the domestic retail side and its welfare effect on domestic households, while directly observing retail prices and quantities, which allow us to measure domestic markups and separate markup adjustments from cost movements. Moreover, the relevant market structure is measured directly from local catchment areas rather than inferred from importer concentration or firm-to-firm sourcing links. This allows us to study a different margin of trade policy and market power: the incidence of a selective tariff exemption that changes consumers’ foreign outside option while leaving domestic retailers’ costs unchanged. The results show that such exemptions can have effects beyond the consumers who use them directly, because they alter the competitive constraint faced by domestic firms.

2 Data

We now describe the battery of administrative datasets, available from 2019 through 2024, used in our analysis. While the datasets are anonymized, variables *across* datasets can be linked via unique (pseudonymous) identifiers at the individual or retailer level.

Customs Data. Our analysis draws on Costa Rican customs microdata. Each import declaration reports the most detailed product classification available (up to

a 10-digit HS code), together with the transaction value, traded quantity, arrival date, and origin country.² As in many settings, the data span firm-level imports. Importantly for our application, they also record shipments imported by *individuals*—for example, a household purchase from an online retailer abroad shipped directly to Costa Rica. The median value of de-minimis individual shipments is about \$21.³ Customs records also include the tariff and taxes paid for each importation.

Price and Product Detail Per Transaction. In Costa Rica, electronic invoices are digital documents used to record sales transactions in compliance with tax regulations. Businesses are required to issue electronic invoices for *all* taxable transactions to simplify tax reporting and reduce evasion. This requirement applies to both firm-to-firm transactions and sales to final consumers, whether in-store or online.⁴

Retailers, in particular, issue one invoice per sale. This invoice includes the retailer’s unique ID, and for a significant share of all sales, it also includes the unique ID of the *final customer who purchased the good*—recall that these IDs are all pseudonymous, but they can be linked across datasets. Not all invoices include this detail, but on average, 42% of transactions do, as it allows for better tracking of transactions and detailed records for both businesses and consumers. In fact, businesses typically encourage customers to provide their ID for each sale, and customers have the added value of keeping an electronic record and invoice of each transaction in their email. For each transaction, these invoices list every product sold, with 13-digit product codes for which, conveniently, the first 10 digits coincide with the Costa Rican HS-10 codes, along with the price and quantity corresponding with each product code. Each invoice also details the retailer’s location.

Household Location and Income. We link each importer or buyer to their residence using official records maintained by the National Registry. While records include district of residence, with 488 districts in total, they also detail the voting center which is closest to each citizen’s residence *for each adult citizen*. With 2,028

²In some product groups an HS-10 classification is not defined, in which case we use the narrowest available HS category (HS-8 or HS-6). Records also include details on the foreign party sending the goods (the shipper/exporter).

³Empirically, nearly all individuals import a given product only once (Argente, Mendez, and Van Patten, 2025), which is inconsistent with systematic stocking behavior. Considering all individual imports, even those above the de-minimis threshold, the median value is \$30.

⁴E-commerce by domestic Costa Rican retailers is, however, limited: online transactions account for less than 3% of retail sales. These sales are nevertheless recorded in the electronic-invoicing system and are therefore included in our analysis.

voting centers in total, the median number of adults assigned to each voting center is only 586.⁵ Therefore, we use the voting center information to construct neighborhoods, leveraging that these assignments provide a standardized and highly granular partition of space. To obtain household income measures, we rely on the Registry of Economic Variables of the Costa Rican Central Bank, which tracks formal employment and labor earnings based on social security records. These employer–employee data include each worker’s occupation and earnings.

2.1 Facts on De Minimis Imports

Before turning to the reform, we briefly characterize the direct-to-consumer import segment affected by our experiment. Namely, throughout this subsection, we restrict attention to individual, non-commercial shipments with declared values at or below the pre-reform de minimis threshold of \$500. First, the typical parcel is small: the median declared value is \$20.61, the mean is \$52.86, and the 90th percentile is \$144.78. Thus, the relevant outside option is not driven by transactions bunching close to the statutory limit, but by many small consumer purchases.

Second, the product mix is highly retail-relevant. The largest HS-4 categories are consumer-facing goods commonly sold by domestic retailers. Women’s and girls’ suits, dresses, skirts, trousers, and shorts (HS-4 6204) account for about 14% of shipments and 11% of value; women’s and girls’ blouses and shirts (HS-4 6206) account for another 11% of shipments and 6% of value. Other leading categories include furniture, luggage and handbags, motor-vehicle parts, toys, knit apparel, footwear, and computers. Table A.1 reports the top categories and their shares of below-threshold de minimis imports. The overlap with domestic retail is economically large: HS-4 categories observed in de minimis shipments account for 65% of the value of domestic retailer sales across HS-4 categories, indicating that the foreign outside option covers a substantial part of the retail product space relevant for our analysis.

Third, origins are also concentrated in countries central to cross-border e-commerce and parcel logistics. The United States accounts for roughly two thirds of below-threshold shipments and about 59% of value, while China accounts for about 22% of shipments and 21% of value. Together, these two origins account for approximately 87% of shipments and 80% of declared value. Details on the main origins are reported in Table A.2.

⁵For details on the distribution of voting centers, see Méndez and Van Patten (2022).

Finally, direct importing is positively correlated with income. Approximately 12% of adults import directly, and consumers who import have 49% higher income than consumers who only buy locally. This pattern is consistent with fixed costs of accessing the de minimis channel: direct importing requires internet access, a debit or credit card, and the ability to navigate foreign websites and overcome language barriers. Thus, judging incidence solely from direct use of the de minimis exemption would point unambiguously toward higher-income households.

3 Background and Natural Experiment

Prior to May 2021, Costa Rica operated a de minimis regime for small direct-to-consumer parcel imports arriving through postal and express-courier channels. Residents could request an exemption from import duties and related charges for non-commercial shipments with a CIF value below US\$500.⁶ The regime created a meaningful duty advantage for imported varieties purchased directly by households, relative to domestically retailed varieties that face full domestic taxation and retail markups.

On May 1, 2021, Costa Rica adopted the Central American Uniform Customs Code and its implementing regulation, in an effort to harmonize customs procedures within Central America. While the overarching objective was administrative standardization and trade facilitation, an incidental implication was the removal of the de minimis exemption channel that consumers had been using for e-commerce purchases under the US\$500 threshold.⁷

After the reform, small parcels purchased online are cleared under the standard import regime. The resulting tax wedge combines (i) the applied import duty, (ii) a 1% levy applied to virtually all imports, and (iii) the domestic value-added tax (13% on most goods), with additional product-specific excises in a limited set of categories. A key institutional feature is that the relevant tariff schedule for these consumer parcels is effectively MFN rather than preferential FTA tariffs. Costa Rica is party to multiple free trade agreements, including CAFTA-DR with the United States, under which most U.S.-origin consumer and industrial goods enter duty free (Méndez and Van Patten, 2022). However, preferential treatment requires the importer to satisfy

⁶The exemption could be claimed by each individual once every six months, leaving a paper trace every time an exception was requested.

⁷After the reform, the duty-free treatment was applied only to the restrictive category of “small family shipments” (*pequeños envíos familiares sin carácter comercial*)—parcels that had to be remitted by family members abroad for the recipient’s personal or family use.

the agreement’s rules of origin—documentation that is routinely provided in brokered firm imports but is costly and rarely furnished by individual consumers clearing one-off parcels. Consistent with this friction, and as the administrative customs data records *actual* tariffs paid by shipment, we observe that duties paid on individuals’ shipments line up with the applied MFN schedule across products.

Importantly, the MFN tariff schedule itself long predates the 2021 reform and was designed around conventional commercial trade flows—intermediate inputs and final goods imported by firms—rather than direct-to-consumer parcel imports. The composition of domestic production in the product space affected by the reform reinforces this interpretation. The leading HS-4 categories imported by individuals under the de minimis regime account for less than 0.2% of Costa Rican manufacturing value added in 2023.⁸ These consumer-facing categories therefore have little domestic manufacturing presence: domestic demand is supplied overwhelmingly through imports rather than local production. Thus, the MFN rates operate as largely dormant tariff shifters that became relevant in our setting only when the reform unexpectedly applied the standard import regime to direct-to-consumer parcels.


In fact, Costa Rica adopted its MFN tariff schedule in January 1995 ([WTO, 1995](#)), which underscores why cross-product variation in MFN rates at the time of the 2021 reform is plausibly orthogonal to the contemporaneous rise of e-commerce parcels and to the political economy of the de minimis regime in Costa Rica. By contrast, commercial imports undertaken by firms (before and after the reform) are typically brokered and documented, and thus claim preferential rates. Therefore, firm-level import tariffs are unaffected by the policy change and equal to zero in virtually all categories, reinforcing the interpretation of the reform as a shock to import competition rather than a domestic-cost shock.

Moreover, the reform likely raised trade costs beyond the average tariff increase. Before May 2021, these parcels were already recorded in customs data and subject to VAT, but they benefited from duty-free treatment. After the reform, the same shipments also required tariff liquidation and payment as part of the release process. This added a fixed shipment-level burden—duty calculation, payment processing, courier collection/handling, and greater exposure to delay or review. Because this administrative component operated through parcel clearance rather than product classification, it was largely common across formerly duty-exempt products.

⁸This calculation excludes the broad plastics category.

Average Tariff Changes The elimination of the de minimis exemption generates a sharp increase in tariffs and taxes paid on affected shipments. Figure 2 illustrates the timing and magnitude of the average change in taxes and tariffs paid for individual shipments below the \$500 threshold, weighted by importing volume, with an average increase of about 8 p.p. (57%).


Figure 2: Policy Change: Elimination of De Minimis Exemption and Tariff Changes



Notes: The figure shows the average tariff, inclusive of taxes, paid before and after the policy reform on non-commercial imports with declared value below \$500, weighted by importing volume. Prior to May 2021, this amount mostly coincided with the VAT. Afterwards, it includes the weighted average tariff.

Variation Across Products. The policy change generates substantial heterogeneity in tariff changes across product categories. Figure 3 visualizes the distribution of tariff changes across HS-4 categories, with changes ranging from zero to 20 p.p.


Figure 3: Variation in Tariff Changes Across HS-4 Categories



Notes: The figure shows the average change in tariff faced by each HS-4 category as a result of the elimination of the de minimis exemption. Changes in tariffs correspond with percentage points.

Variation Across Neighborhoods. Because neighborhoods differ in pre-policy importing patterns across product categories, the elimination of the de minimis rule generates substantial heterogeneity in the average tariff change faced across space. Figure 4 shows cross-neighborhood variation in import-weighted tariff changes.


Figure 4: Variation in Average Tariff Change Across Neighborhoods



Notes: The figure shows the average tariff change across neighborhoods, weighted by import value, resulting from the elimination of the de minimis exemption. Changes in tariffs are measured in percentage points.

Consumer Import Response. The policy change generated a large, plausibly exogenous increase in the delivered price of households' foreign outside option and led to a sharp decline in direct-to-consumer importing. Figure 5 shows that the decline occurred immediately after the reform. Panel (a) shows a roughly 65% decrease in the number of unique importers per quarter, while Panel (b) shows a roughly 40%-50% decline in the value of imports by individuals. As discussed above, the reform combined the removal of duty-free treatment with shipment-level clearance costs, helping rationalize why the observed contraction is larger than what would be implied by the average 8 percentage-point tariff increase alone.

Figure 5: Drop in Consumer Imports Following the Policy Change



Notes: The figure shows the decline in consumer imports following the policy change, restricting attention to non-commercial individual imports with a declared value under \$500. Panel (a) considers the total number of unique importers per month, and Panel (b) shows the total value imported by individuals per quarter.

A Competition Shock. Finally, it is worth highlighting that a key feature of this reform is that, in principle, it changes the *price of the foreign option* available to households while leaving domestic retailers' costs unaffected—which we will confirm empirically. Thus, it provides a direct, product-specific shock to the competitiveness of local retailers.

4 Catchment Areas and De Minimis Exposure

4.1 Estimating Retailer-Specific Catchment Areas

Retailers serve different sets of neighborhoods. We construct *retailer-specific catchment areas* using the geographic distribution of customers observed in electronic invoices: a retailer's catchment area is the collection of neighborhoods from which its customers originate. This approach is feasible because a sizable share of retail invoices record the buyer's national ID. In our data, 42% of retail electronic invoices include a customer ID, on average across years. Table B.2 summarizes the distribution of the share of transactions with customer IDs among retailers.

Consumers frequently provide their ID because it is convenient: the electronic invoice is automatically delivered to the email registered with the tax authority, allowing households to keep receipts for their own records, while retailers value customer IDs to track repeat customers and often incentivize their provision. Consistent with

this, 96% of adults in our sample appear at least once as an identified buyer in a retail invoice, and every retailer in our sample records at least some transactions with identified customers.

We operationalize catchment areas using the distribution of a retailer’s sales across customers’ neighborhoods of residence. For each retailer, we compute the share of sales in the customer-linked transaction sample that is accounted for by buyers residing in each neighborhood. Intuitively, if two thirds of retailer r ’s sales among transactions that can be linked to a neighborhood are to customers from neighborhood A , and one third are to customers from neighborhood B , then the retailer’s catchment area places weight $2/3$ on A and $1/3$ on B , with weights summing to one across neighborhoods. This sales-weighted construction is especially useful in our setting because it avoids the coarse approximations often used in retail studies, such as administrative districts or a fixed-distance radius around store locations. It also naturally accommodates disconnected catchment areas: for example, a store located in a commercial district may primarily serve customers who reside in multiple residential neighborhoods.

More formally, we define retailer r ’s catchment area as the empirical distribution of its sales across customers’ neighborhoods of residence. Let Sales_{rh} denote retailer r ’s sales to customers residing in neighborhood h , measured using transactions for which the buyer can be linked to a neighborhood. We define the weight of neighborhood h in retailer r ’s catchment area as

$$w_{rh} \equiv \frac{\text{Sales}_{rh}}{\sum_j \text{Sales}_{rj}}. \quad (1)$$

The catchment area of retailer r is then the support of this distribution,


$$Z_r \equiv \{h : w_{rh} > 0\}. \quad (2)$$

That is, a retailer’s catchment area is not defined mechanically, but by the empirical distribution of the neighborhoods from which its customers are drawn. This definition will map naturally into our exposure measure, which aggregates neighborhood-level import exposure using the same weights.

Figure 6 illustrates retailer-specific catchment areas as sets of neighborhoods h belonging to the support of retailer r ’s customer base. Following Equations (1) and (2), each colored neighborhood is assigned a weight w_{rh} , given by its share of retailer r ’s sales. The color gradients represent these catchment weights, with darker shades

corresponding to neighborhoods that account for a larger share of the retailer's sales. The figure also highlights that catchment areas need not be mutually exclusive; as an example, neighborhoods shown in purple belong to the catchment areas of both retailers.

Figure 6: Examples of Retailer Catchment Areas




Notes: The figure presents an illustrative example of retailer-specific catchment areas for two retailers. Each polygon represents a neighborhood, and retailers are denoted by storefront icons. Colored neighborhoods belong to the support of at least one retailer's catchment area. Following equations (1) and (2), each neighborhood h receives a weight $w_{r,h}$ equal to its share of retailer r 's sales. The color gradients represent these weights, with darker shades indicating neighborhoods that account for a larger share of the retailer's sales. Turquoise and red denote the catchment areas of the two retailers, while purple denotes neighborhoods that belong to both catchment areas.

Implications. The estimated catchment areas reveal a retail environment that is spatially segmented and locally concentrated. Table B.1 reports the distribution of the number of neighborhoods in each retailer's catchment area. Many retailers serve geographically narrow customer bases: the median retailer draws customers from only three neighborhoods, and retailers at the 25th percentile are highly local: they draw customers from a single neighborhood. At the same time, the distribution has a long upper tail, with retailers at the 90th percentile serving customers from 56 neigh-

borhoods. This heterogeneity underscores why fixed-radius or administrative-market definitions would poorly approximate the relevant customer base across retailers.

To better understand concentration and market power, Figure 7 plots the distribution of the maximum retailer market share within neighborhood–HS-4 markets. Concentration varies widely across markets, and a substantial mass lies in the upper tail, where a single retailer accounts for a large fraction of local sales. In fact, over 12% of neighborhood–HS-4 markets have one retailer accounting for all observed local retail sales. This upper-tail concentration is central to the mechanism studied below: when the foreign outside option becomes more expensive, these are the markets in which weaker import competition can potentially translate into large markup and price responses.

Figure 7: Distribution of Maximum Market Shares Across Neighborhood–HS-4s



Notes: The figure shows the distribution of the maximum retailer market share across neighborhood–HS-4 markets. Market shares are computed using pre-policy domestic retail sales in each local product market. Product categories are defined at the HS-4 level.

A Note on Customer IDs. Appendix B.2 formalizes the implications of observing customer IDs for only a subset of retail invoices, deriving how sampling error in the estimated catchment weights propagates into our outcomes of interest and relating this issue to the granular-uncertainty framework of Dingel and Tintelnot (2026). In our setting, three features limit the scope for this issue. First, customer IDs are not used to measure prices, quantities, product codes, retailer identifiers, or total sales, all of which are observed for the universe of electronic invoices; IDs are used only to estimate w_{rh} . Second, these weights are constructed from many identified customer

links per retailer, as shown in Table B.2, so the relevant uncertainty is sampling error in a catchment-area average rather than sparsity in a high-dimensional retailer–neighborhood matrix. Third, Figure B.1 shows that ID coverage is not systematically related to the key objects in the analysis: retailer market shares, catchment-area income, or catchment-area exposure. Thus, any residual error in $\hat{w}_{rh}$ would primarily attenuate measured exposure and market-share variation. It would not generate the upper-tail price responses documented below unless ID disclosure were systematically correlated with exposure and market power in precisely the way ruled out by the coverage diagnostics.

Catchment-Area Timing and Alternative Constructions. In the baseline analysis, we hold catchment areas fixed rather than allowing the relevant market definition to change mechanically before and after the reform. Specifically, we construct each retailer’s catchment area using post-reform shopping links and then use this fixed geography to measure exposure, market shares, and outcome responses. The motivation is that the reform raises the cost of the foreign option and may therefore reveal a broader set of domestic shopping alternatives. Post-reform catchments thus provide a conservative, maximal approximation to each retailer’s latent customer base once direct importing becomes less attractive. The construction is deliberately inclusive: a neighborhood enters a retailer’s catchment area if at least one resident purchased any product from that retailer in the two years following the reform. We then study how exposure reallocates expenditure and changes prices within this fixed potential market.

Appendix B.3 provides two complementary validation exercises for this catchment-area construction. First, using retailers with high customer-ID coverage as a benchmark sample, we estimate a PPML gravity model that predicts catchment weights from bilateral distance, neighborhood characteristics, and retailer characteristics. We then use the predicted weights to reconstruct w_{rh} for all retailers. Second, we construct an alternative set of catchment weights using the *pre-policy* residence neighborhoods of each retailer’s employees as a proxy for the geography of its local footprint. For both alternative constructions, we recompute the key exposure and market-share objects and re-estimate the main outcomes. In both exercises, the results align closely with the baseline construction based on customers’ residences; moreover, the baseline estimates lie between the two alternative estimates.

The three constructions use distinct sources of information and therefore discipline

complementary aspects of the market definition. Customer-residence catchments recover the geography of realized shopping links; the PPML reconstruction replaces raw links with predicted spatial weights based on geography and observables; and employee-residence catchments proxy for the retailer’s local footprint without using customer purchases. The similarity of the results across these approaches indicates that the main findings do not hinge on a particular feature of the customer-based measure or on measurement error specific to any single construction. These exercises are also useful beyond our setting: the PPML reconstruction speaks to environments in which bilateral distance and retailer–neighborhood characteristics are observed but customer links are not, while the employee-residence approach provides a second proxy based on administrative employer–employee records, which are often available even when customer locations are not.

4.2 Retailer-Product Exposure

The joint availability of product codes in electronic invoices and HS codes in customs records allows us to connect the foreign outside option for a product group to local retail outcomes. This mapping is especially clean because product codes in Costa Rica’s electronic invoicing system are built directly on the Harmonized System: the first 10 digits of the 13-digit invoice code coincide exactly with the HS-10 classification reported in customs records. As a result, imports and retail transactions can be linked directly. In the empirical analysis and model calibration, we define product groups as HS-4 categories.

Let g denote a product group (HS-4), r a retailer, and h a neighborhood. Let $S_{h,fg0}$ denote the pre-policy share of direct-to-consumer imports in local consumption of category g for neighborhood h . Because retailer r serves multiple neighborhoods, we summarize the foreign competitive pressure it faces in category g by taking a catchment-area weighted average of these neighborhood-level foreign shares:

$$\bar{S}_{fg0}^r \equiv \sum_h w_{rh} S_{h,fg0}, \quad \text{where} \quad w_{rh} \equiv \frac{\text{Sales}_{rh}}{\sum_j \text{Sales}_{rj}}, \quad (3)$$

where w_{rh} is the share of retailer r ’s sales accounted for by neighborhood h . We then define retailer-product exposure to the reform as

$$IC_{rg} \equiv \bar{S}_{fg0}^r \Delta\tau_g, \quad (4)$$

where $\Delta\tau_g$ is the policy-induced change in the tariff-and-tax burden on consumer imports in category g , expressed in percentage points.⁹ Intuitively, IC_{rg} is larger when a retailer sells disproportionately to neighborhoods that relied heavily on de minimis imports in category g before the reform, and when category g experiences a larger increase in the cost of importing under the new regime.

A key advantage of this measure is that the cross-product component, $\Delta\tau_g$, is plausibly exogenous to contemporaneous retail outcomes. As discussed above, the reform replaced the previous zero-duty treatment for eligible consumer parcels with the pre-existing MFN schedule. These MFN rates were set decades before the reform and with commercial trade by firms—including intermediates and final goods—in mind, rather than direct-to-consumer parcel imports. In this sense, the reform delivers a sharp and predetermined shift in the foreign outside option, while the retailer-specific component of exposure is governed by pre-policy importing patterns across the neighborhoods in each retailer’s catchment area.

5 Empirical Strategy

We estimate horizon-specific long-difference specifications that compare a common pre-reform baseline to outcomes measured at different horizons after the reform. Let $T \in \mathcal{T}$ index a post-policy horizon. For any outcome z , define

$$\Delta_T \ln z_i \equiv \ln z_{i,T} - \ln z_{i,0},$$

where $z_{i,0}$ is the average of outcome z in the pre-policy period and $z_{i,T}$ is the corresponding average in the post-policy window associated with horizon T . Estimating separate long differences by horizon allows the effects of the reform to evolve flexibly over time, while abstracting from high-frequency noise, transitory fluctuations, and short-run adjustment dynamics.

Our baseline specifications are estimated at the retailer-product level. Let $y_{rg,t}$ denote an outcome for retailer r in product category g and period t , where g indexes HS-4 categories. Depending on the specification, $y_{rg,t}$ is either the average retail price or total sales of retailer r in category g . For each horizon T , we estimate

⁹Thus, IC_{rg} is measured in percentage-point exposure units: for example, if a category’s tariff-and-tax burden rises by 7 percentage points and the baseline foreign share in retailer r ’s catchment area is 0.10, then $IC_{rg} = 0.7$.

$$\Delta_T \ln y_{rg} = \beta_T IC_{rg} + \alpha_T \bar{S}_{fg0}^r + \lambda_r + \lambda_g + \varepsilon_{rgT}, \quad (5)$$

where IC_{rg} is the retailer-product exposure measure defined in [equation \(4\)](#).¹⁰ Furthermore, $\bar{S}_{fg0}^r$ denotes the baseline exposure share underlying the shift-share construction, λ_r denotes retailer fixed effects, λ_g denotes product-category fixed effects, and ε_{rgT} is an error term, clustered by product-category. Including $\bar{S}_{fg0}^r$ together with retailer and product-category fixed effects implies that β_T is identified from variation in IC_{rg} that is orthogonal to the baseline share structure and fixed effects. This isolates the component of the shift-share exposure that is not driven by cross-sectional differences in initial shares, in the spirit of the residualized (or “recentered”) variation emphasized in [Borusyak, Caceres-Bravo, and Hull \(2025\)](#). Moreover, $\bar{S}_{fg0}^r$ absorbs any component of the reform that is common across formerly exempt products but proportional to a retailer–category cell’s pre-policy reliance on the foreign option.¹¹

The two-way fixed effects play an important role in isolating the mechanism of interest. Since IC_{rg} combines a retailer-specific component (the pre-policy foreign share in the retailer’s catchment area) with a category-specific component (the tariff change), the retailer fixed effects λ_r absorb any retailer-level change common across categories, such as changes in local demand, store-wide pricing policies, management, sourcing conditions, or retailer-specific cost shocks over the horizon considered. Likewise, the category fixed effects λ_g absorb any category-wide change common across retailers, such as national demand shifts, changes in product-level costs, or secular trends in category g .

As a result, the coefficient β_T is identified from the interaction of category-level tariff shocks with predetermined cross-retailer differences in pre-policy exposure. Namely, as the reform is nationwide, identification does not come from a comparison between treated and untreated retailers. Instead, it comes from differences in *exposure intensity*. In the full-sample specification, and given the retailer and product fixed effects, the comparison group for a highly exposed retailer-product cell (r, g) is

¹⁰Because IC_{rg} is measured in percentage-point exposure units and the dependent variables are log changes, the coefficients β_T are semi-elasticities: they measure log-point changes in the outcome associated with a one-percentage-point increase in exposure.

¹¹This is relevant because, as discussed above, the reform may have raised the cost of direct importing not only through statutory tariffs but also through common parcel-clearance frictions. Thus, the coefficient on IC_{rg} is identified from the product-specific tariff component of the reform, net of common changes associated with baseline foreign reliance.

a set of less-exposed cells: on the one hand, other product categories sold by the same retailer that were less exposed to the reform; on the other hand, the same product category sold by other retailers whose catchment areas were less reliant on de minimis imports before the policy. Put differently, β_T is identified by whether retailer-product cells with higher pre-policy exposure experience larger post-reform changes than less-exposed cells, after netting out retailer-wide and category-wide changes.

For outcomes defined at the retailer level, such as employment or average wages, we aggregate exposure across product categories using pre-policy sales weights.¹² We also construct the retailer-level analogue of the baseline exposure-share control, $\bar{S}_{f0}^r$. As retailer-level outcomes cannot be estimated with the same two-way fixed-effect structure as the retailer-product specifications in [equation \(5\)](#), we condition on a battery of retailer controls and fixed effects that absorb differences in product mix, activity, baseline reliance on foreign competition, and scale adjustment. For each horizon T , we estimate

$$\Delta_T \ln x_r = \gamma_T IC_r + \mathbf{X}_{rT}' \rho_T + \eta_{g(r)} + \eta_{n(r)} + \eta_{p(r)} + u_{rT}, \quad (7)$$

where x_r is either total employment or average wages. The vector of retailer controls includes $\bar{S}_{f0}^r$, and average tariff and total revenue changes faced by the retailer. The fixed effects $\eta_{g(r)}$, $\eta_{n(r)}$, and $\eta_{p(r)}$ absorb the retailer's main HS-4 category, neighborhood, and industry classification, respectively. The coefficient γ_T therefore compares retailers with similar product mix, activity, baseline exposure shares, and scale adjustment, but different policy-induced exposure IC_r .

6 Reduced Form Results

We now present reduced-form evidence on how the elimination of the de minimis exemption affected domestic retail outcomes. The central result is that retailer-HS-4 cells more exposed to the reform experienced larger increases in both prices and sales. Crucially, these price increases are not accompanied by corresponding increases

¹²Specifically, define retailer-level exposure as

$$IC_r \equiv \sum_g \omega_{rg} IC_{rg}, \quad \text{where} \quad \omega_{rg} \equiv \frac{\text{Sales}_{rg,0}}{\sum_k \text{Sales}_{rk,0}}, \quad (6)$$

and where $\text{Sales}_{rg,0}$ denotes retailer r 's pre-policy sales in category g , so that ω_{rg} is the share of retailer r 's baseline sales accounted for by category g .

in domestic retailers’ cost proxies: exposed retailer–HS-4 cells show no meaningful rise in import prices or domestic purchase prices, and more exposed retailers do not significantly expand employment or raise wages once we condition on product-mix, activity, and scale controls. This pattern supports the mechanism developed above. The reform raises the delivered price of households’ foreign outside option, shifts residual demand toward domestic retailers, and leaves domestic retailers’ marginal costs largely unchanged. The resulting increase in retail prices is therefore consistent with a markup response to weaker import competition rather than cost pass-through.

Prices and Sales. Table 1 reports the baseline two-year long-difference estimates at the retailer–HS-4 level. Retailer–product cells with greater exposure to the reform experience larger increases in both prices and sales. A one-percentage-point increase in IC_{rg} raises prices by approximately 2.6 percent, and sales by approximately 1.8 percent.¹³ Thus, in categories where the policy increased the cost of direct importing more strongly, domestic retailers charge more and record higher nominal sales.

The relative magnitudes of the price and sales responses are informative. The positive sales coefficient shows that exposed domestic retailer–product cells receive higher nominal expenditure after the reform. However, nominal sales combine two forces: expenditure reallocated from foreign direct-to-consumer sellers toward domestic retailers, and higher prices induced by weaker foreign competition. Because the specification controls for the baseline foreign share $\bar{S}_{fg0}^r$, it compares retailer–product cells with similar ex-ante scope for reallocation but different tariff-induced increases in the cost of the foreign option. Along this margin, the price response exceeds the nominal-sales response, implying a decline in implied quantities. This pattern is consistent with the market-power mechanism: once direct foreign purchases become more expensive, domestic retailers gain residual market power and raise markups; the resulting price increase dominates the quantity expansion that would arise from substitution toward domestic sellers.

Ruling Out Cost Pass-Through. Table 2 shows that the reform changed the competitiveness of foreign direct-to-consumer imports without materially altering the cost structure of domestic retailers. Institutionally, this is to be expected: the policy changes the treatment of parcels imported by households, but it does not change the tariff schedule applied to retailers’ standard commercial sourcing. We are able

¹³The mean value of exposure IC_{rg} is 0.7 percentage points.

Table 1: Reduced-Form Impact on Retail Prices and Sales

	$\Delta \ln(\text{Price}_{rg})$ (1)	$\Delta \ln(\text{Sales}_{rg})$ (2)
IC_{rg}	0.026*** (0.005)	0.018*** (0.004)
Retailer FE	✓	✓
HS-4 FE	✓	✓
$\bar{S}_{fg0}^r$	✓	✓
Observations	321,322	321,322

Notes: Standard errors in parentheses. The table reports two-year long-difference estimates at the retailer–HS-4 level. The regressions include retailer fixed effects, HS-4 fixed effects, and the baseline foreign-share control $\bar{S}_{fg0}^r$. The exposure variable IC_{rg} is measured in percentage-point exposure units. Because dependent variables are log changes, coefficients are semi-elasticities. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

to study this directly, as customs data include retailers’ import prices by product for firms that import directly, while electronic invoice data record product prices in firm-to-firm transactions for retailers purchasing from other local firms.

Consistent with the institutional setting, columns (1) and (2) show no meaningful increase in retailer import prices or purchase prices in affected product categories. The import-price coefficient and the purchase-price coefficient are both small relative to the retail price response in Table 1 and statistically indistinguishable from zero. This evidence speaks against a cost-pass-through interpretation. If retail prices were rising because retailers were paying more to source goods, we would expect corresponding increases in import or purchase prices. We do not observe such increases. The result also weighs against an interpretation in which higher retail prices reflect post-reform quality upgrading, as such upgrading would be expected to raise retailers’ purchase prices.

Columns (3) and (4) provide a complementary check using retailer-level labor outcomes. They show that more exposed retailers do not exhibit statistically significant increases in employment or wages. Thus, the price response is not accompanied by evidence of higher labor costs, hiring, or a broader increase in operating costs.

Taken together, results in Table 1 and Table 2 indicate that the reform does not appear to increase retailers’ input costs, labor costs, quality upgrading, or overall operating costs. Yet, exposed retailers still raise final prices. This pattern is consistent with a shift in residual demand toward domestic retailers that allows them to charge higher markups, supporting our interpretation of the reform as a shock to effective market power.

Table 2: Input and Labor Costs

	Input prices		Labor outcomes	
	$\Delta \ln(\text{Import Price}_{rg})$ (1)	$\Delta \ln(\text{Purchase Price}_{rg})$ (2)	$\Delta \ln(\text{Employment}_r)$ (3)	$\Delta \ln(\text{Wage}_r)$ (4)
IC_{rg}	0.0006 (0.016)	0.003 (0.003)		
IC_r			0.004 (0.006)	0.005 (0.004)
-----	-----	-----	-----	-----
Retailer FE	✓	✓		
HS-4 FE	✓	✓		
$\bar{S}_{fg0}^r$	✓	✓		
Retailer controls $\mathbf{X}_{rT}$			✓	✓
Main HS-4 FE			✓	✓
Neighborhood FE			✓	✓
Industry FE			✓	✓
Observations	7,767	163,104	20,752	20,752

Notes: Columns (1)–(2) are retailer–HS-4 regressions. Column (1) uses the unit import price, $\log(\text{dollar}_M/\text{quantity}_M)$, for retailers that import directly. Column (2) uses the purchase-price outcome from firm-to-firm input purchases. These regressions include retailer and HS-4 fixed effects and control for the baseline foreign share $\bar{S}_{fg0}^r$. Columns (3)–(4) are retailer-level labor regressions using aggregated exposure IC_r . The vector of retailer controls is $\mathbf{X}_{rT} = (\bar{S}_{fg0}^r, \bar{\Delta\tau}_r, \Delta_T \ln R_r)'$, where $\bar{S}_{fg0}^r$ is the retailer-level baseline foreign-share control, $\bar{\Delta\tau}_r$ is the sales-weighted tariff-change component, and $\Delta_T \ln R_r$ is the change in total retailer revenues. These regressions also include fixed effects for the retailer’s main HS-4 category, neighborhood, and industry classification. Standard errors in parentheses. Standard errors are clustered by HS-4 in columns (1)–(2) and by retailer in columns (3)–(4). Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Assortment, Entry/Exit, and Substitution. The reform also affects the set of varieties available to households. Table 3 reports two complementary pieces of evidence, chosen to mirror the nesting structure that we will build into the model in the next section. Columns (1)–(2) ask whether the composition of options within affected HS-4 categories changes after the reform, considering the average number of fine varieties sold by a retailer within an HS-4 category g . Column (1) focuses on domestic retailers and shows that, in more exposed categories, domestic retailers expand their assortments, while Column (2) shows that varieties sold by foreigners contract. Thus, once direct importing becomes more expensive, domestic retailers expand their presence in the affected product groups, while foreign direct-to-consumer supply contracts. These assortment responses show that the reform changes the effective set of options available within HS-4 categories, not just the prices of a fixed basket of goods.

Column (3) asks whether the contraction of the foreign option leads households to reallocate expenditure away from the affected HS-4 categories altogether. For each neighborhood–HS-4 cell, we construct the change in the share of total neighborhood expenditure allocated to category g , Δs_{hg} , and regress it on a neighborhood–category

exposure measure IC_{hg} , including neighborhood and HS-4 fixed effects.¹⁴ The coefficient is economically and statistically insignificant, indicating that affected categories do not lose expenditure share within neighborhoods. Thus, households that had been buying a category from abroad do not simply stop buying that category or shift expenditure to other HS-4s. Instead, the evidence points to substitution within the same product category, from foreign direct-to-consumer sellers toward domestic retailers.

Table 3: Assortment Responses and Expenditure Reallocation

	Assortment		Substitution
	$\Delta \text{Varieties}_g^{\text{Sold by Domestic}}$ (1)	$\Delta \text{Varieties}_g^{\text{Sold by Foreign}}$ (2)	Δs_{hg} (3)
IC_g	1.124** (0.462)	-0.862*** (0.211)	
IC_{hg}			-0.00009 (0.002)
$\bar{S}_{fg0}$	✓	✓	
HS-4 controls $\mathbf{X}'_{g0}$	✓	✓	
$\bar{S}_{fg0}^h$ FE			✓
HS-4 FE			✓
Neighborhood FE			✓

Notes: Columns (1)–(2) are HS-4-level regressions. In column (1), the outcome is the change in the average number of fine varieties sold by domestic retailers within HS-4 category g . The coefficient is measured in varieties per domestic retailer. Column (2) reports the corresponding change in foreign direct-import varieties within HS-4 category g . The variable IC_g denotes category-level exposure to the reform. The vector of HS-4 controls $\mathbf{X}_{g0}$ includes HS-4-level baseline expenditure, the sales-weighted tariff-change component, and baseline domestic and foreign variety counts. Column (3) is estimated at the neighborhood–HS-4 level. The outcome s_{hg} is neighborhood h 's expenditure share allocated to HS-4 category g , and $IC_{hg} = S_{h,fg0} \Delta \tau_g$ is the neighborhood–category exposure measure. Column (3) includes neighborhood and HS-4 fixed effects and controls for the baseline foreign share $S_{h,fg0}$. Standard errors are in parentheses. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Reduced Form Robustness. Tables B.4 and B.5 show that the baseline results, which use customer residence neighborhoods from invoices to construct catchment areas, are robust to two alternative constructions: a PPML gravity reconstruction of catchment-area weights and a proxy based on employees' residence neighborhoods. In both cases, the estimates align closely with the baseline, with the baseline magnitudes lying between the two alternatives. This indicates that the main results are not

¹⁴Specifically, $s_{hg,t} \equiv E_{hg,t} / \sum_k E_{hk,t}$, where $E_{hg,t}$ is total expenditure by residents of neighborhood h in HS-4 category g . The exposure measure used in this column is $IC_{hg} \equiv S_{h,fg0} \Delta \tau_g$, where $S_{h,fg0}$ is the pre-policy direct-import share in neighborhood h 's expenditure on HS-4 g . Thus, IC_{hg} is the local analogue of the exposure measure used in the retailer-level analysis, but applied to the upper-nest budget-share outcome.

driven by the particular source of geographic information used to construct retailer catchment areas.

Moreover, the baseline estimates below use a two-year long difference. Appendix C.1 reports analogous and consistent estimates at a one-year horizon. As an additional product-definition robustness check, Appendix C.2 repeats these regressions using HS-6 product categories instead of HS-4 categories, both at a one- and a two-year horizon. The price and sales responses remain positive at the narrower product level, while input-cost measures do not increase.

7 A Model of Variable Markups with Retailer-Specific Catchment Areas

This section develops a quantitative framework that maps the policy-induced increase in the cost of direct-to-consumer importing into retail price changes and welfare effects. The model is built to mirror three features emphasized in the reduced-form analysis: (i) direct imports provide a foreign outside option at the product-category level; (ii) retailers set prices for their catchment areas rather than for a single national market; and (iii) there is a markup response to weakened import competition.

The model is closely related to the variable-markup environment in Atkeson and Burstein (2008), but departs from the standard Atkeson-Burstein setup as market shares are not defined in an economy-wide sectoral market. Instead, they are defined over *retailer-specific catchment-area markets* constructed from household shopping links. This wedge allows the same foreign cost shock to generate very different markup and price responses across locations, even within the same product group.

7.1 Preferences

Neighborhoods are indexed by h and time by t . Households have CES preferences over a continuum of product groups and CES preferences within product groups across

retailer options:

$$U_{ht} = \left[\int_{g \in \Omega_h} \phi_{hgt} C_{hgt}^{\frac{\gamma-1}{\gamma}} dg \right]^{\frac{\gamma}{\gamma-1}}, \quad (8)$$

$$C_{hgt} = \left[\sum_{r \in \Omega_{hgt}} \phi_{hrgt} C_{hrgt}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad 1 < \gamma < \sigma. \quad (9)$$

Households face the budget constraint $P_{ht}C_{ht} = E_{ht}$. The set Ω_{hgt} includes a foreign retailer option that represents direct consumer importing.

The nesting is chosen to mirror the substitution margin we observe in the data. Column (3) of Table 3 shows that the reform does not meaningfully change the share of expenditure allocated to affected HS-4 categories. Thus, the main adjustment occurs *within* HS-4 categories: consumers substitute from the foreign option toward domestic retailers, rather than away from the product category altogether. This motivates an upper nest over HS-4 product groups and a lower nest over retailer options within each HS-4.

Three aspects of the aggregation are worth emphasizing. First, the upper-level aggregator in equation (8) treats product groups as a continuum, delivering the standard CES allocation of expenditure across HS-4 categories. Second, and crucially for market power, the within-group aggregator in equation (9) sums over a finite set of retailer options available to households in neighborhood h for product group g . This discreteness allows individual retailers to hold non-negligible expenditure shares in local neighborhood–product markets.¹⁵ This feature is empirically important in our setting where local concentration is substantial, as shown in Figure 7, with a thick upper tail of neighborhood–HS-4 markets in which a single retailer accounts for a large share of sales. Third, it is also useful to note that each retailer–HS-4 pair represents a composite variety. In our implementation, this composite is constructed by aggregating (log) prices over finer product codes using expenditure-share weights, which corresponds to an implicit Cobb–Douglas lower layer.¹⁶

¹⁵If retailer options instead formed a continuum within g , each individual retailer’s share would be infinitesimal and the perceived demand elasticity would collapse to the constant σ . In that limiting case, a shock that only raises the price of the foreign option would reallocate expenditure away from foreign purchases but would not generate the markup-driven domestic price increases we document, since domestic marginal costs are unchanged.

¹⁶Adding a CES nest within retailer–HS-4 cells would require estimating an additional elasticity of substitution across finer products. It would not, however, change the strategic-complementarity channel, which operates through the retailer–HS-4 catchment-area market share S_{rgt} .

The nested CES structure in equations [equation \(8\)](#) and [equation \(9\)](#) delivers a demand elasticity faced by each retailer that is *endogenously* decreasing in its market share. This demand curvature is what allows retail prices to respond to the policy even when domestic marginal costs remain unchanged: when the foreign outside option becomes more expensive, retailers gain share, the perceived elasticity falls, and markups rise. Importantly, because market shares are defined over retailer-specific catchment areas in the implementation below, the same foreign-price shock generates heterogeneous markup responses across retailers even within the same HS-4 category.¹⁷

7.2 Retail Pricing and Catchment-Area Market Shares

Retailer r sells product category g to households in the neighborhoods that make up its catchment area Z_r . We assume that the retailer sets a uniform price P_{rgt} for category g across these neighborhoods and faces constant marginal cost A_{rgt} . For single-establishment retailers—92% of firms in our data—uniform pricing is not a discretionary modeling assumption. It is the relevant empirical constraint: the store cannot condition the storefront price on the customer’s residential neighborhood. For multi-establishment firms, the common-price restriction is also consistent with existing evidence that retail prices are typically uniform, or nearly uniform, across locations.¹⁸ For all firms, price is chosen against the residual demand generated by the particular neighborhoods they serve. Taking competitors’ prices and neighborhood expenditures as given, retailer r ’s problem is given by

¹⁷A natural extension would be to allow for non-homothetic CES demand. Introducing non-homotheticity in the within-category retailer nest would require estimating income-specific substitution parameters across retailer–category varieties. A more tractable alternative would be to place non-homotheticity only in the upper nest, so that income shifts expenditure across product categories. This shifter affects demand levels but drops out of the log pricing first-order condition, so the residual demand elasticity faced by the retailer, and hence the estimated markup-elasticity schedule, is unchanged conditional on the retailer’s market share. The drawback is in welfare accounting: once price indices depend on income, compensating variation no longer decomposes cleanly into a common cost-of-living component and an income component. We therefore use the homothetic benchmark to isolate the price-index channel that is directly disciplined by observed price and market-share responses.

¹⁸[DellaVigna and Gentzkow \(2019\)](#) show that large U.S. food, drugstore, and mass-merchandise chains charge nearly uniform prices across stores despite large differences in local demographics and competition, with related evidence for online and multichannel retailers ([Nakamura, 2008](#); [Cavallo, 2018](#); [Daruich and Kozlowski, 2023](#)).

$$\begin{aligned}
& \max_{P_{rgt}} \quad \Pi_{rgt} = (P_{rgt} - A_{rgt}) Y_{rgt}, \\
& \text{subject to} \quad Y_{rgt} = \sum_{h \in Z_r} Y_{rgt}^h,
\end{aligned} \tag{10}$$

where Y_{rgt}^h denotes demand for retailer r 's variety of category g from households in neighborhood h . This formulation mirrors the empirical construction: a retailer's total sales in a product category are the collection of the sales it makes across the neighborhoods it serves.

This pooled-demand structure makes incidence intrinsically cross-neighborhood. The price paid by households in a given neighborhood h is not determined only by demand and competitive conditions in h . Instead, P_{rgt} is chosen to maximize profits over the entire catchment area, so the relevant residual demand elasticity depends on conditions across all neighborhoods served by the retailer. As a result, a shock that raises the price of the foreign outside option for households in neighborhood h' can affect the price paid by households in neighborhood h whenever both neighborhoods belong to the same retailer's catchment area. These interdependencies further justify defining the relevant market at the retailer's catchment-area level, rather than treating each neighborhood in isolation. They also highlight why reduced-form exposure alone may miss the equilibrium effect of the tariff change: a shock to the foreign option in one part of a retailer's customer base can alter the retailer's optimal price, and therefore affect consumers throughout the catchment area.

The first order condition implies the following.¹⁹

$$P_{rgt} = \frac{\sigma(1 - S_{rgt}) + \gamma S_{rgt}}{\sigma(1 - S_{rgt}) + \gamma S_{rgt} - 1} A_{rgt}, \tag{11}$$

where

$$S_{rgt} \equiv \sum_{h \in Z_r} \omega_{hrgt} S_{hrgt}, \quad \omega_{hrgt} \equiv \frac{E_{hrgt}}{\sum_{\ell \in Z_r} E_{\ell rg t}}. \tag{12}$$

Equation (11) has the classic structure of price equaling a markup term multiplied by the marginal cost. The markup, in turn, depends on S_{rgt} , which is the retailer's *catchment-area market share*. This share is an expenditure-weighted average of its local market shares across the neighborhoods it serves (equation (12)). This is the key

¹⁹See Appendix F.1 for the step-by-step derivation.

departure from the standard [Atkeson and Burstein \(2008\)](#) environment. In Atkeson–Burstein, a firm’s market share is defined within an integrated sectoral market. Here, market shares are retailer specific because each retailer serves a different set of neighborhoods. Two retailers selling the same HS-4 category need not compete over the same set of consumers, and therefore need not face the same pricing incentives. In the empirical implementation, Z_r and the weights entering S_{rgt} are disciplined by the shopping links in electronic invoices, as described in [Section 4.1](#).

An Exact Mapping from Shares to Prices. A key advantage of our setting is that we observe catchment-area market shares S_{rgt} . This allows us to connect price changes to share changes using the model’s exact pricing relationship rather than relying only on a first-order approximation. Define the retailer’s perceived elasticity and markup as functions of its catchment-area market share:

$$\varepsilon_{rgt} \equiv \sigma(1 - S_{rgt}) + \gamma S_{rgt}, \quad \mu_{rgt} \equiv \frac{\varepsilon_{rgt}}{\varepsilon_{rgt} - 1}.$$

Then [equation \(11\)](#) can be written as

$$\log P_{rgt} = \log A_{rgt} + \log \mu_{rgt}.$$

Holding domestic marginal costs fixed, the change in domestic prices is therefore

$$\Delta \log P_{rgt} = \log \mu_{rgt} - \log \mu_{rg,t-1}. \quad (13)$$

Thus, observed changes in catchment-area market shares can be translated directly into implied changes in markups and prices for any candidate pair (σ, γ) . It is also useful to express this mapping in differential form; namely,

$$d \log P_{rgt} = d \log \mu_{rgt} = \frac{\Gamma_{rgt}}{\sigma - 1} d \log S_{rgt}. \quad (14)$$

where

$$\Gamma_{rgt} \equiv (\sigma - 1) \frac{\partial \log \mu_{rgt}}{\partial \log S_{rgt}} = \frac{(\sigma - \gamma)(\sigma - 1)S_{rgt}}{\varepsilon_{rgt}(\varepsilon_{rgt} - 1)}. \quad (15)$$

[Equation \(15\)](#) is the key object that we discipline empirically. It depends on the retailer–product-specific catchment-area market share S_{rgt} , rather than on a national or sectoral market share. This departure from the standard [Atkeson and Burstein \(2008\)](#) environment implies that two retailers selling the same HS-4 category can

have different markup elasticities because they serve different catchment areas and therefore have different local market shares. The object also depends on the two nest elasticities (σ, γ) : σ governs substitution across retailer options within an HS-4 category, while γ governs substitution across product groups.

Exposure as a Demand Shifter (and an Instrument). In empirical applications, market shares are often unobserved, which leads to using a first-order approximation to map the foreign price shock into domestic price changes. In our context, if the foreign price rises by

$$d \log P_{gt}^f \approx \frac{\tau_{gt} - \tau_{g,t-1}}{1 + \tau_{g,t-1}},$$

then a standard first-order argument implies

$$d \log P_{rgt} \approx \underbrace{\Gamma_{rgt} \left[S_{fg,t-1}^r \left(\frac{\tau_{gt} - \tau_{g,t-1}}{1 + \tau_{g,t-1}} \right) \right]}_{\text{Exposure}}, \quad (16)$$

where $S_{fgt}^r \equiv \sum_{h \in Z_r} \omega_{hrgt} S_{hfgt}$ is the foreign share in the retailer's catchment area.

In our setting, besides observing price changes at the retailer-product level, we observe (i) the foreign-share component that governs the strength of the demand shifter, and (ii) the domestic market shares that determine the markup elasticity. We therefore do not need to treat exposure as a reduced-form proxy for price changes, nor do we need to rely on instrumenting the leave-out price index.²⁰

Instead, we use exposure as an *instrument* for the *observed* change in the retailer–HS-4's catchment-area market share, and then discipline (σ, γ) using the differential relationship in [equation \(14\)](#). This is the key step that, in the next section, allows us to recover the implied markup elasticities across the market-share distribution, disciplined by quasi-experimental variation in import competition.

²⁰For instance, under strategic complementarities, [Amiti et al. \(2019a\)](#) rely on the relationship that $d \log P_{rgt} = \frac{\Gamma_{rgt}}{1 + \Gamma_{rgt}} d \log \mathcal{P}_{rgt}^{-r}$ where $\Gamma \equiv -\frac{\partial \log \mu_{rgt}}{\partial \log P_{rgt}}$. We do not do so, both because of the observed shares but also as the leave-out price index is a much more complicated object given the catchment-area structure:

$$d \log \mathcal{P}_{rgt}^{-r} \equiv \frac{\sum_{h \in Z_{rgt}} \omega_{rgt}^h S_{rgt}^h \sum_{j \in \Omega_{gt}^h, j \neq r} S_{jgt}^h d \log P_{jgt}}{\sum_{h \in Z_{rgt}} \omega_{rgt}^h S_{rgt}^h (1 - S_{rgt}^h)}.$$

7.3 Model Calibration And Fit

We discipline the two key demand elasticities (γ, σ) using the policy-induced shifts in competition and the observed joint behavior of prices and market shares across the share distribution. This calibration strategy targets the object most directly relevant for incidence. Rather than asking the model to fit a broad set of quantity or expenditure moments, (γ, σ) are disciplined using the empirical price–share elasticity schedule: how $\Delta \log P_{rg}$ responds to quasi-experimental changes in $\Delta \log S_{rg}$ at different points of the market-share distribution. This is the key object governing how a change in the foreign outside option maps into local markup and price responses. Estimating the parameters from this targeted heterogeneity makes the resulting incidence calculations less dependent on the model’s ability to match other margins that are not central to the market-power channel.

The identifying observation is that the model’s strategic complementarity term Γ_{rgt} in [equation \(15\)](#) depends *only* on the market share S_{rgt} and the two elasticities (γ, σ) . Since we observe market shares, we can recover the mapping between market structure and markup elasticities from the data, following four steps.

Step 1: Shares, Prices, and Exposure per Catchment Area. The model-consistent catchment-area shares S_{rgt} are the same catchment areas constructed from electronic invoices in [Section 4.1](#). Retail prices P_{rgt} are measured as average transaction prices at the retailer-HS-4 level. Finally, our predetermined exposure shifter is the same shift-share exposure used in the reduced-form exercises, IC_{rg} .

Step 2: Use Exposure to Instrument for Changes in Market Shares. [Equation \(14\)](#) relates changes in prices to changes in market shares, and through [equation \(15\)](#) can allow us to recover the key elasticities. However, even though we observe both prices and shares, share changes are not mechanically exogenous: they can also reflect local demand shocks, retailer-specific shocks, or category trends. We therefore instrument these market shares using exposure, IC_{rg} . As Γ depends on S_{rgt} , it is natural to do this estimation in bins of the market share distribution. For bin b , let $\mathcal{B}_b$ denote the set of retailer–category cells in that bin. The estimated 2SLS coefficient would be given by

$$\widehat{\kappa}_b^{IV} \equiv \frac{\widehat{\beta}_b^p}{\widehat{\beta}_b^s}, \quad \Delta y_{rg} = \beta_b^y IC_{rg} + \alpha_b^y \bar{S}_{fg0}^r + \lambda_r^y + \lambda_g^y + \varepsilon_{rg}^y, \quad y \in \{p, s\}, (r, g) \in \mathcal{B}_b, \quad (17)$$

where $\Delta p_{rg} \equiv \Delta \log P_{rg}$ and $\Delta s_{rg} \equiv \Delta \log S_{rg}$. Thus, $\hat{\kappa}_b^{IV}$ measures the causal elasticity of prices with respect to market shares within share bin b , driven by the foreign competition shock.²¹

Step 3: Recover Elasticities (γ, σ) . The model implies that the price-share elasticity in [equation \(17\)](#) is pinned down by (γ, σ) through

$$\kappa^{\text{model}}(S; \gamma, \sigma) = \frac{\Gamma(S; \gamma, \sigma)}{\sigma - 1}, \quad \Gamma(S; \gamma, \sigma) = \frac{(\sigma - \gamma)(\sigma - 1)S}{\varepsilon(S)(\varepsilon(S) - 1)}, \quad \varepsilon(S) = \sigma(1 - S) + \gamma S. \quad (18)$$

Because $\kappa^{\text{model}}(\cdot)$ depends on the market share S and only two unknown parameters, matching the empirical elasticity at two points in the share distribution yields two equations in two unknowns. Concretely, letting $S^{(1)}$ and $S^{(2)}$ denote representative shares (e.g., the average share within a bin), we solve

$$\hat{\kappa}(S^{(1)}) = \kappa^{\text{model}}(S^{(1)}; \gamma, \sigma), \quad \hat{\kappa}(S^{(2)}) = \kappa^{\text{model}}(S^{(2)}; \gamma, \sigma),$$

for (γ, σ) . Intuitively, the lower-bin share pins down the baseline curvature, while the higher-bin disciplines the steepening of markup elasticities in the upper tail, which is the central nonlinearity in the data. We choose these two bins to minimize the distance between the empirical and model-based distributions of the elasticity of markups.


This procedure yields $\sigma = 7$ and $\gamma = 1.5$. The implied within-group elasticity σ is close to the median across-brand elasticity reported in [Broda and Weinstein \(2010\)](#), consistent with relatively strong substitution across retailer options within a product category. The across-group elasticity γ governs substitution across broad product groups and is lower, implying more limited substitution across categories.²²

Step 4: Fit Across the Full Share Distribution. Having identified (γ, σ) , we compute $\kappa^{\text{model}}(S; \gamma, \sigma)$ for all values of S and plot it against the empirical $\hat{\kappa}_b$ across market-share bins. [Figure 8](#) shows that the model replicates the key empirical feature: the elasticity is modest at low shares and rises sharply in the upper tail.

²¹The first-stage share equation and the reduced-form price equation are estimated with the same controls and fixed effects as in [equation \(5\)](#).

²²This hierarchy is consistent with the reduced-form evidence in [Table 3](#), which shows little reallocation of neighborhood expenditure shares across affected HS-4 categories.

Figure 8: Markup Elasticities Increase with Retailer Market Share



Notes: The figure shows the relationship between retailer market share and the elasticity of markups. The empirical series reports estimates from the 2SLS specification. The model series reports the corresponding model-implied elasticities, given the estimated elasticities from the structural model.

Intuitively, once the reform raises the delivered price of the foreign option, retailers in low-market-share environments remain disciplined by other local competitors. By contrast, in high-share markets, the foreign option is often one of the few remaining competitive constraints. Once that constraint weakens, the retailer gains substantial market power and can sharply increase prices. These high-share markets are also quantitatively important. As shown in Figure G.1, the top decile of the market-share distribution accounts for 18% of domestic sales.²³

7.4 Robustness Exercises

Results Based on Alternative Catchment Area Constructions. Appendix B.3 shows that the estimated markup-elasticity schedule is robust to alternative catchment-area constructions. Figure B.2 repeats the IV exercise after recomputing catchment weights using two alternatives: a PPML gravity reconstruction of retailer-neighborhood links and a proxy based on employees' residences. Across the market-

²³For completeness, Figure G.2 shows that input-price responses do not display the same systematic upper-tail pattern as retail prices. Input-cost estimates are close to zero over the market-share distribution, supporting the interpretation that the steep upper-tail retail price response reflects markups rather than cost pass-through.

share distribution, the resulting schedules closely track the baseline customer-based measure. All three constructions imply modest markup elasticities at low and intermediate shares and the same sharp steepening in the upper tail. Moreover, the baseline schedule lies between, and very close to, the two alternative schedules.

Results Based on National Market Shares. Appendix D then asks what would happen if the same exercise used national retailer shares within each product category. Figure D.1 shows that the resulting markup-elasticity schedule is essentially flat and informative only at very low market shares. The reason is that, as Table D.1 shows, national retailer–HS-4 shares are very small: the 90th percentile is only 0.1%, and even the 99th percentile is 3%. One might expect national shares to perform better when product markets are defined more narrowly. We therefore use national–HS-6 shares in Figure D.2. The schedule remains essentially flat, and the distribution of national shares remains concentrated near zero. As shown in Table D.2, the 90th percentile of national–HS-6 shares is only 0.2%, and even the 99th percentile is 7%. Thus, even at a narrower product definition, national shares contain little variation in the upper-tail region.

By contrast, the local catchment-area measures recover this high-share variation and reveal a sharply increasing schedule in the upper tail. Together with the robustness in Figure B.2, this comparison highlights the value of measuring market power at the local level. Even when catchments are constructed from less direct *but locally grounded* information—PPML-predicted links or employee residences—they preserve the high-share variation that is central to the mechanism. National shares, by contrast, aggregate across heterogeneous pricing markets and therefore wash out the upper-tail variation where the theory predicts markup elasticities should be largest.

Results for Single-Establishment Retailers. We also repeat the main exercises restricting attention to single-establishment retailers. Table E.1 shows that the reduced-form estimates are statistically equal to the baseline estimates. More exposed single-establishment retailer–HS-4 cells raise both prices and sales, while input-cost measures do not increase. Figure E.1 estimates the markup-elasticity schedule in the same subsample and shows the same central pattern as in the full sample; if anything, the upper-tail response is somewhat steeper for single-establishment firms, consistent with the idea that multi-establishment retailers with common pricing may pool demand conditions across broader pricing markets and thereby attenuate purely local price responses. Overall, the exercise strengthens the robustness of the baseline re-

sults: the reduced-form effects and the nonlinear markup-elasticity schedule are not driven by multi-establishments or by assumptions about chain-level pricing, but are already present among firms tied to a single storefront.

8 Welfare and Distributional Effects

We quantify welfare impacts using compensating variation at the neighborhood level. The key advantage of the model-based approach is that it allows us to translate heterogeneous retailer-level price and assortment responses into heterogeneous changes in neighborhood cost of living, and then into welfare effects that can be compared across the income distribution.

Compensating Variation. Let P_t^h denote the neighborhood price index and I_t^h total income. For small changes, compensating variation (as a fraction of initial expenditure) can be written as

$$\widetilde{CV}_t^h \approx \% \Delta P_t^h - \% \Delta I_t^h, \quad (19)$$

so welfare losses arise from higher cost of living and lower income. In our setting, the reduced-form results suggest that income changes through wages and employment are limited, implying that heterogeneity in welfare effects is primarily driven by heterogeneity in price index changes and by changes in profits and variety.

Cost of Living. We construct the neighborhood price index using observed price changes among continuing retailer options and an adjustment for entry and exit:

$$\% \Delta P_h = \left[\prod_{g \in \Omega_h} (EPI_g^h)^{\omega_g^h} \right] - 1, \quad EPI_g^h = \left(\frac{\lambda_{hgt}}{\lambda_{hg,t-1}} \right)^{\frac{1}{\sigma-1}} \prod_{r \in \Omega_g^h} \left(\frac{P_{hrgt}}{P_{hrg,t-1}} \right)^{\omega_{rg}^h}, \quad (20)$$

where λ_{hgt} is the share of expenditure on continuing stores and ω are log-expenditure weights. This structure speaks to the reduced-form assortment evidence. The entry/exit term $(\lambda_{hgt}/\lambda_{hg,t-1})^{1/(\sigma-1)}$ captures the welfare-relevant effect of changes in the set of available retailer options. Thus, the extensive-margin changes documented earlier enter the cost-of-living index through the Feenstra variety-adjustment term, both through the domestic variety expansion and the foreign variety contraction. The

heterogeneous price changes documented in the reduced form map instead into the intensive-margin price-index component.

Income Changes. Neighborhood income changes reflect labor income, profits, and tariff rebates:

$$\% \Delta I_h = \omega_L^h \% \Delta W_h + \omega_\Pi^h \% \Delta \Pi_h + \omega_T^h \% \Delta T_h. \quad (21)$$

Empirically, we find no changes in wages or employment, in Table 2. However, as the tariff weakens the foreign outside option, it raises domestic retailer markups, and shifts rents toward domestic retailers and the government. Whether those rents compensate households depends on how profits and tariff revenue are redistributed.

Because the incidence of domestic retail profits and tariff revenues is not directly observed, our baseline welfare calculation holds nominal income fixed and focuses on cost-of-living changes. We then report alternative calculations that rebate incremental domestic retail profits across neighborhoods under explicit incidence assumptions. Namely, these rents are allocated across neighborhoods in proportion to the additional retail expenditure generated by residents of each neighborhood.

8.1 Results and Heterogeneity by Market Structure

We now aggregate the model-implied price-index changes to the neighborhood level and evaluate welfare incidence using compensating variation. Our baseline calculation holds nominal income fixed, thereby isolating the cost-of-living effect of the reform. We also report a redistribution-inclusive exercise that allocates incremental domestic retail profits across neighborhoods under an explicit incidence assumption.

The model implies modest median welfare effects but substantial heterogeneity across neighborhoods. In the fixed-income baseline, with compensating variation expressed as a percentage of initial expenditure, the median compensating variation is 0.03%, while the interquartile range is 0.09 percentage points.


Figure 9 plots neighborhood-level compensating variation against the expenditure-weighted maximum market share faced by residents.²⁴ A positive compensating variation corresponds to a welfare loss: the percentage of initial expenditure households would need to be compensated to offset the reform-induced increase in the cost of

²⁴Market shares are computed to capture local market power. Local market power is measured as the expenditure-weighted average of the maximum retailer market share across neighborhood-HS-4 markets, using residents' baseline expenditure weights. Thus, the horizontal axis captures the markup-relevant market share faced by the average dollar of neighborhood expenditure.

living. Panel (a) holds nominal income fixed and focuses on cost-of-living changes. Panel (b) allocates incremental retail profits across neighborhoods in proportion to the additional domestic retail expenditure generated by residents of each neighborhood. In both cases, welfare losses are largest in high-share, low-competition neighborhoods. Holding nominal income fixed, households experience welfare losses across the full market-power distribution. Including the retail-profit rebate generates negative compensating variation in some low-market-power neighborhoods. This occurs because, under the rebate assumption, part of the additional domestic retail profits generated by substitution away from the foreign option is allocated back to residents of those neighborhoods. These gains are largest where substitution toward domestic retailers is substantial but markup-driven price increases remain limited.

This pattern summarizes the paper’s central distributional implication. The welfare consequences of eliminating the de minimis exemption are not determined only by which households imported directly before the reform, but also by the local market structure they face when they substitute toward domestic retailers. When the foreign outside option becomes more expensive, households in concentrated markets face larger price increases and have limited scope to reallocate toward competing local sellers, generating larger welfare losses.

Figure 9: Compensating Variation Results




Notes: The figure plots neighborhood-level compensating variation against local market power. Local market power is measured as the expenditure-weighted average of the maximum retailer market share across neighborhood-HS-4 markets, using residents’ baseline expenditure weights. Compensating variation is expressed as a percentage of initial expenditure; for example, a value of 2 corresponds to 2% of initial expenditure. Under our sign convention, positive values indicate welfare losses. In Panel (a), fixed nominal income implies that compensating variation equals the change in the neighborhood price index. Panel (b) allocates incremental retail profits across neighborhoods in proportion to the additional domestic retail expenditure generated by residents of each neighborhood.

Figure 10 relates the main objects driving incidence to *income*, measured by the average wage of households. The figure highlights two opposing forces. Panel (a)

shows that exposure to the de minimis reform is increasing in income: higher-income areas relied more heavily on direct-to-consumer imports before the policy change. Judged only through this exposure margin, the reform would therefore appear progressive. Panel (b), however, shows that local market power moves in the opposite direction: retailer market shares are highest in lower-income areas, where households face fewer effective domestic alternatives. Together with Figure 9, these patterns imply that losses are largest in lower-income areas, despite their lower direct exposure to de minimis imports. Thus, once endogenous markup responses and local concentration are taken into account, a policy that appears progressive based only on who imported directly can be regressive in equilibrium.

Figure 10: Exposure vs. Market Share by Income Level



Notes: The figures plot neighborhood–HS-4 level outcomes against average wage income. Panel (a) reports exposure to the de minimis reform. Panel (b) reports local market power, measured as the expenditure-weighted average of the maximum retailer market share across neighborhood–HS-4 markets, using residents’ baseline expenditure weights. Thus, Panel (b) captures the markup-relevant market share faced by the average dollar of neighborhood expenditure. The figure shows that exposure is highest in higher-income neighborhoods, while local market power is largest in lower-income neighborhoods.

The costs of remoteness are also evident.²⁵ Beyond income, the drivers of incidence vary systematically with an area’s remoteness, measured by its distance to the nearest international gateway (port or airport). Figure G.3 shows that policy exposure is highest in locations closer to an international gateway, where access to imported goods is likely easiest. By contrast, market shares are larger in more remote areas, consistent with weaker local competition farther from those gateways. As in Figure 10, these two channels work in opposite directions. This pattern is not mechanical in the Costa Rican context: although San José and many higher-income neighborhoods

²⁵For related terminology, see Redding and Sturm (2008) on the costs of remoteness; see also Allen et al. (2022) on the “curse” of remoteness.

are located in the Central Valley, the country’s major ports are on the coasts, and only one international airport is near the capital. Geographic remoteness therefore shapes incidence in ways that are distinct from the income gradient.

9 Conclusion

De minimis regimes have become a central institution of modern trade. By lowering the all-in cost of buying abroad, they expand households’ choice sets and strengthen the foreign outside option that disciplines domestic retailers.²⁶ This paper studies what happens when that outside option is abruptly made more expensive. We exploit Costa Rica’s May 2021 elimination of duty-free treatment for non-commercial parcels below \$500 as a large and plausibly exogenous increase in the delivered price of direct-to-consumer imports.

Using linked administrative data that combine shipment-level customs records for individual imports with the universe of retail electronic invoices and household locations and earnings, we construct retailer-specific catchment areas and measure retailer-product exposure as the interaction of product-level tariff changes with pre-reform foreign import penetration in each catchment area. In the data, more exposed retailer-HS-4 cells raise both prices and nominal sales, while retailers’ import prices and purchase prices do not rise and labor outcomes (employment and wages) do not meaningfully increase. These patterns are difficult to reconcile with cost pass-through and instead point to an increase in domestic markups driven by weakened import competition. At the same time, assortment adjusts: exposed categories gain domestic retail varieties and lose foreign varieties, indicating that the reform reshapes both intensive- and extensive-margin components of consumer welfare.

A central finding is that incidence depends sharply on local market structure. Retail price responses are highly nonlinear in pre-policy market shares: they are modest in low-share markets but steep in the upper tail, where concentration is highest. To interpret these facts and quantify welfare effects, we develop a catchment-area model of spatial competition with variable markups and strategic complementarities. Relative to the standard Atkeson-Burstein framework, the key departure is that market shares, and therefore markups, are defined over *retailer-specific catchment-area mar-*

²⁶While de minimis shipments are small relative to total trade by value, they constitute a large share of cross-border e-commerce shipments by count. For the United States, the USITC reports that Section 321 (de minimis) imports are a substantial share of e-commerce imports by quantity and estimates that by value they were less than 2 percent of total U.S. goods import value in 2021.

kets rather than integrated sectoral markets. A distinctive feature of our setting is that we observe the relevant domestic and foreign shares directly, allowing us to use the policy-induced foreign price shock to instrument for observed share changes and discipline the markup-elasticity schedule. The model matches the steep upper-tail nonlinearity and implies modest median welfare effects but substantial heterogeneity: welfare losses are concentrated in low-competition, lower-income neighborhoods.

These findings speak directly to current policy debates over *de minimis*. For example, in the United States, the White House issued an Executive Order on July 30, 2025 suspending duty-free *de minimis* treatment under 19 U.S.C. 1321(a)(2)(C) (Section 321 of the Tariff Act of 1930)—historically a channel for low-value shipments to enter without duties—with implementation effective August 29, 2025 and subsequently continued in 2026.²⁷ This episode illustrates that tightening *de minimis* is no longer a hypothetical reform but an active instrument of trade policy.

More broadly, the mechanism we document applies *beyond* *de minimis*. Any policy that raises the delivered price of imports—including tariffs, non-tariff barriers, and enforcement frictions that disproportionately burden small shipments—can relax the foreign competitive constraint and induce endogenous increases in domestic markups. In particular, the results show that the incidence of selective tariff exemptions extends beyond the consumers who use them directly. Because local retail competition is spatially segmented and highly heterogeneous, the consumer incidence of tariffs need not track exposure alone; it can be amplified in concentrated markets and thus generate regressive outcomes even when the directly exposed importers are relatively high income. Incorporating local market structure and variable markups is therefore essential for evaluating the consumer and distributional consequences of trade policy.

Beyond international trade, the same mechanism speaks to the distributional impacts of online retail. Online shopping can discipline local market power by expanding households’ outside options, especially in places where retail markets are concentrated and consumers have few nearby alternatives. Conversely, policies or frictions that raise the cost of accessing online sellers can relax this competitive constraint and increase local markups. The extent of these effects depends on how retailers price across space: common pricing by chains may pool shocks across neighborhoods and dampen some local differences, while more segmented pricing can make the gains from online access

²⁷The implementation notice specifies the suspension of the duty-free *de minimis* exemption for covered products valued at \$800 or less, with operational details and exceptions for certain flows.

sharply uneven. Thus, the incidence of online retail depends not only on who buys online, but also on how online options reshape competition in local retail markets.

References

- Rodrigo Adão, Ana M. Fernandes, Chang-Tai Hsieh, and Jose M. Quintero. Trade shocks in distorted economies: Evidence from firm-level import data. Working paper, version dated June 6, 2025, June 2025.
- Treb Allen, David Atkin, Santiago Cantillo, and Daniel Hernández. Trucks or the triple curse of remoteness. Working paper, 2022.
- Vanessa I. Alviarez, Michele Fioretti, Ken Kikkawa, and Monica Morlacco. Two-sided market power in firm-to-firm trade. NBER Working Paper 31253, National Bureau of Economic Research, 2023.
- Mary Amiti, Oleg Itskhoki, and Jozef Konings. Importers, exporters, and exchange rate disconnect. *American Economic Review*, 104(7):1942–1978, 2014. doi: 10.1257/aer.104.7.1942.
- Mary Amiti, Oleg Itskhoki, and Jozef Konings. International shocks, variable markups, and domestic prices. *The Review of Economic Studies*, 86(6):2356–2402, 2019a. doi: 10.1093/restud/rdz005.
- Mary Amiti, Stephen J. Redding, and David E. Weinstein. The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives*, 33(4):187–210, 2019b. doi: 10.1257/jep.33.4.187.
- David Argente, Esteban Mendez, and Diana Van Patten. Cross-border product adoption: Individual imports, migrant networks, and domestic retailers, 2025. Working paper.
- Costas Arkolakis, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. The elusive pro-competitive effects of trade. *Review of Economic Studies*, 86(1): 46–80, 2019. doi: 10.1093/restud/rdx075.
- Andrew Atkeson and Ariel Burstein. Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031, 2008. doi: 10.1257/aer.98.5.1998.
- David Atkin, Benjamin Faber, and Marco Gonzalez-Navarro. Retail globalization and household welfare: Evidence from Mexico. *Journal of Political Economy*, 126(1): 1–73, 2018. doi: 10.1086/695476.
- Alex Blumenfeld, Christina Patterson, and Joseph Vavra. Concentrating on customers: A new measure of size-based market power. Work in progress, presentation slides, May 2026.

- Kirill Borusyak, Maria Caceres-Bravo, and Peter Hull. Estimating demand with recentered instruments. Working paper, 2025.
- Christian Broda and David E Weinstein. Globalization and the Gains from Variety. *The Quarterly Journal of Economics*, 121(2):541–585, 2006.
- Christian Broda and David E Weinstein. Product creation and destruction: Evidence and price implications. *American Economic Review*, 100(3):691–723, 2010.
- Alberto Cavallo. More amazon effects: Online competition and pricing behaviors. Working Paper 25138, National Bureau of Economic Research, October 2018.
- Alberto Cavallo, Gita Gopinath, Brent Neiman, and Jenny Tang. Tariff pass-through at the border and at the store: Evidence from U.S. trade policy. *American Economic Review: Insights*, 3(1):19–34, March 2021. doi: 10.1257/aeri.20190536.
- Diego Daruich and Julian Kozlowski. Macroeconomic implications of uniform pricing. *American Economic Journal: Macroeconomics*, 15(3):64–108, 2023. doi: 10.1257/mac.20210172.
- Jan De Loecker, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. Prices, markups, and trade reform. *Econometrica*, 84(2):445–510, March 2016. doi: 10.3982/ECTA11042.
- Stefano DellaVigna and Matthew Gentzkow. Uniform pricing in U.S. retail chains. *The Quarterly Journal of Economics*, 134(4):2011–2084, 2019. doi: 10.1093/qje/qjz019.
- Jonathan I. Dingel and Felix Tintelnot. Spatial economics for granular settings. *Econometrica*, 94(2):407–464, 2026. doi: 10.3982/ECTA19350.
- Te Du, Chao Fang, Shuzhong Ma, Lixin Tang, and Zhihong Yu. Moving online under de minimis? import tariffs and international retail e-commerce. Technical report, CESifo Working Paper, 2025.
- Pablo D. Fajgelbaum and Amit Khandelwal. The value of De Minimis imports. Working Paper 32607, National Bureau of Economic Research, June 2024. URL <https://www.nber.org/papers/w32607>. Revision date: February 2025.
- Pablo D. Fajgelbaum and Amit K. Khandelwal. Measuring the unequal gains from trade. *The Quarterly Journal of Economics*, 131(3):1113–1180, August 2016. doi: 10.1093/qje/qjw013.
- Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. The return to protectionism. *The Quarterly Journal of Economics*, 135(1): 1–55, February 2020. doi: 10.1093/qje/qjz036.
- Robert C. Feenstra. New product varieties and the measurement of international prices. *American Economic Review*, 84(1):157–177, March 1994. doi: 10.2307/2117976. URL <https://www.jstor.org/stable/2117976>.

- Robert C. Feenstra and David E. Weinstein. Globalization, markups, and US welfare. *Journal of Political Economy*, 125(4):1040–1074, 2017. doi: 10.1086/692695.
- Pinelopi Koujianou Goldberg and Nina Pavcnik. Distributional effects of globalization in developing countries. *Journal of Economic Literature*, 45(1):39–82, March 2007. doi: 10.1257/jel.45.1.39.
- Paul R. Krugman. Increasing returns, monopolistic competition, and international trade. *Journal of International Economics*, 9(4):469–479, November 1979. doi: 10.1016/0022-1996(79)90017-5.
- Marc J. Melitz and Gianmarco I. P. Ottaviano. Market size, trade, and productivity. *The Review of Economic Studies*, 75(1):295–316, January 2008. doi: 10.1111/j.1467-937X.2007.00463.x.
- Esteban Méndez and Diana Van Patten. Voting on a trade agreement: Firm networks and attitudes toward openness. Technical report, National Bureau of Economic Research, 2022.
- Simon Mongey and Michael E. Waugh. Pricing inequality. Working Paper 33399, National Bureau of Economic Research, January 2025. URL <https://www.nber.org/papers/w33399>.
- Monica Morlacco. Market power in input markets: Theory and evidence from french manufacturing. Unpublished manuscript, 2019.
- Emi Nakamura. Pass-through in retail and wholesale. *American Economic Review*, 98(2):430–437, 2008. doi: 10.1257/aer.98.2.430.
- Guido G. Porto. Using survey data to assess the distributional effects of trade policy. *Journal of International Economics*, 70(1):140–160, September 2006. doi: 10.1016/j.jinteco.2005.09.003.
- Stephen J. Redding and Daniel M. Sturm. The Costs of Remoteness: Evidence from German Division and Reunification. *American Economic Review*, 98(5):1766–1797, 2008. doi: 10.1257/aer.98.5.1766.
- The White House. Suspending duty-free De Minimis treatment for all countries, July 2025. URL <https://www.whitehouse.gov/presidential-actions/2025/07/suspending-duty-free-de-minimis-treatment-for-all-countries/>. Executive Order, July 30, 2025.
- The White House. Continuing the suspension of duty-free De Minimis treatment for all countries. Executive Order, February 2026. URL <https://www.whitehouse.gov/presidential-actions/2026/02/continuing-the-suspension-of-duty-free-de-minimis-treatment-for-all-countries/>. February 20, 2026.
- Petia Topalova. Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India. *American Economic Journal: Applied Economics*, 2(4):1–41, October 2010. doi: 10.1257/app.2.4.1.

WTO. Trade Policy Review: Costa Rica (1995). WTO Trade Policy Review highlights page, 1995. URL https://www.wto.org/english/tratop_e/tpr_e/tp007_e.htm. Highlights page noting that Costa Rica adopted a new tariff schedule in January 1995; accessed 2026-02-27.

Online Appendix for

When Tariffs Raise Markups: Import Competition and Local
Market Power

June 29th, 2026

David Argente
Yale University and NBER

Esteban Méndez
Central Bank of Costa Rica

Diana Van Patten
Yale University and NBER

Table of Contents

Appendix A. De Minimis Import Facts	1
A.1 Additional Tables	1
A.2 Gravity and De Minimis Imports	2
Appendix B. Details on Customer IDs Across Retailers	3
B.1 Statistics on Identified Transactions per Retailer	3
B.2 Identified Customers and Granular Uncertainty	5
B.3 Alternative Catchment-Area Constructions	6
B.4 Results with Alternative Catchment-Area Constructions	8
Appendix C. Reduced Form: Additional Robustness	11
C.1 Alternative Horizon	11
C.2 Results by HS-6 Category	12
Appendix D. National Market Shares by Product Category	13
D.1 National Shares by HS-4	13
D.2 National Shares by HS-6	14
Appendix E. Results for Single-Establishment Retailers	15
Appendix F. Model Details	16
F.1 Derivation of the Pricing Rule	16
Appendix G. Model-Relevant Empirical Regularities	17

A De Minimis Import Facts

A.1 Additional Tables

Table A.1: Top HS-4 Categories in De Minimis Imports

HS-4	Short description	Shipment share	Value share
6204	Women's/girls' suits, dresses, skirts, trousers, shorts	13.6%	10.7%
6206	Women's/girls' blouses, shirts, and shirt-blouses	10.6%	5.8%
9403	Other furniture and parts thereof	1.5%	2.8%
4202	Trunks, suitcases, handbags, wallets, and similar containers	4.1%	2.6%
8708	Parts and accessories of motor vehicles	1.0%	2.5%
9503	Toys, dolls, puzzles, and similar articles	2.2%	1.8%
6110	Sweaters, pullovers, sweatshirts, and similar knit articles	1.5%	1.6%
3926	Other articles of plastics	1.9%	1.6%
6402	Footwear, rubber/plastic soles and uppers	1.5%	1.6%
8471	Computers and automatic data-processing units	0.4%	1.5%

Notes: The table reports the largest HS-4 categories among individual, non-commercial imports with declared values at or below \$500. Shipment shares are computed using the number of below-threshold de minimis parcels. Value shares are computed using declared import value. Product descriptions are shortened HS descriptions.

Table A.2: Top Origins of De Minimis Imports

Shipments		Value	
Origin	Share	Origin	Share
United States	65.6%	United States	58.7%
China	21.9%	China	21.1%
Colombia	1.4%	Mexico	2.5%
Vietnam	1.3%	Colombia	2.3%
Mexico	1.2%	Panama	2.0%
Panama	1.1%	Vietnam	1.0%
India	0.5%	Spain	0.8%
Cambodia	0.4%	Hong Kong	0.8%
Indonesia	0.4%	Peru	0.7%
Spain	0.4%	Canada	0.6%

Notes: The table reports the largest origin countries among individual, non-commercial imports with declared values at or below \$500. The left panel ranks origins by shipment counts, while the right panel ranks origins by declared import value. Shipment shares are computed using the number of below-threshold de minimis parcels. Value shares are computed using declared import value.

A.2 Gravity and De Minimis Imports

This appendix documents that de minimis imports remain shaped by standard trade frictions, even though they are direct-to-consumer and often mediated by online platforms and express logistics. We estimate a PPML gravity specification at the origin-year level using the value of individual, non-commercial imports with declared values at or below \$500:

$$\mathbb{E}[M_{ot} \mid X_{ot}] = \exp(\beta_1 \log GDP_{ot} + \beta_2 \log Dist_o + \beta_3 Cont_o + \beta_4 Lang_o + \beta_5 Col_o + \eta_t),$$

where M_{ot} is the value of below-threshold de minimis imports from origin o in year t , GDP_{ot} is origin GDP, $Dist_o$ is bilateral distance to Costa Rica, and η_t denotes year fixed effects. Standard errors are clustered by origin.

Table A.3: Gravity and De Minimis Imports

	De minimis imports ($\leq \$500$)
Origin GDP	1.748*** (0.100)
Distance	-1.174*** (0.121)
Contiguity	2.755*** (0.529)
Common language	0.415 (0.481)
Colonial dependency	0.454 (0.432)
Year FE	Yes
Pseudo R^2	0.957
Observations	382

Notes: The table reports PPML estimates of [equation \(A.2\)](#). The dependent variable is the value of individual, non-commercial imports from origin o in year t , restricted to shipments with declared values at or below \$500. Origin GDP and distance enter in logs. Standard errors, clustered by origin, are in parentheses. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The estimates show that the below-threshold de minimis channel is not frictionless. Imports are strongly increasing in origin market size and decreasing in distance, and contiguity has a large positive association with import values. Common language and colonial ties are positive but not precisely estimated. Thus, even for small direct-to-consumer parcels, the foreign outside option depends on the geography and logistics of international trade.

B Details on Customer IDs Across Retailers

B.1 Statistics on Identified Transactions per Retailer

Table B.1: Distribution of Retailer Catchment Areas

	Percentile				
	10th	25th	50th	75th	90th
Neighborhoods in a retailer's catchment area	1	1	3	13	56


Notes: The table reports the distribution across retailers of the number of neighborhoods in each retailer's estimated catchment area, where a neighborhood belongs to retailer r 's catchment area if $w_{rh} > 0$.

Table B.2: Share of Transactions with Customer IDs: Distribution Across Retailers

	Percentile					Moments	
	10th	25th	50th	75th	90th	Mean	Std. Dev.
Share of transactions with customer ID	0.01	0.04	0.23	0.91	1.00	0.42	0.39
Mean number of IDs per retailer	139	189	983	2,010	3,290	—	—

Notes: The table reports the distribution across retailers of the share of final-consumer transactions that include a customer ID. The first row reports selected percentiles, the mean, and the standard deviation of the retailer-level ID coverage rate. The second row reports the mean number of observed customer IDs per retailer within each percentile bin. Transactions with firms as counterparties are excluded from the calculation, since these transactions are legally required to include an ID.

Figure B.1: Share of Transactions with Customer ID by Retailer Across Outcomes



Notes: The figures show the share of transactions with customer IDs across different distributions, by retailer. Panels (a), (b), and (c) consider these transactions against market share by retailer, and wages and exposure by catchment area, respectively.

B.2 Identified Customers and Granular Uncertainty

A natural concern is that retailer catchment areas are constructed from the subset of invoices for which the buyer’s ID is reported, rather than from buyer IDs for the universe of retail transactions. This concern is conceptually distinct from missing outcome data. Prices, quantities, product codes, seller identifiers, and total sales are observed for the universe of electronic invoices; the buyer ID is used only to assign a subset of final-consumer sales to residence neighborhoods and thereby estimate the catchment weights w_{rh} . Thus, the relevant question is whether the estimated weights $\hat{w}_{rh}$ are sufficiently precise and representative.

Following the logic of [Dingel and Tintelnot \(2026\)](#), granularity becomes problematic when very sparse realized links in a high-dimensional matrix are treated as if they were precisely estimated probabilities. Our empirical object is less demanding. We do not rely on a handful of observations in each individual retailer–neighborhood cell. Instead, for each retailer r , we use identified invoices to estimate retailer-level averages,

$$x_r \equiv \sum_h w_{rh} x_h,$$

where, to construct catchment areas, x_h are retailer sales. Low-probability neighborhoods receive low weight by construction, and the sampling error in this average is governed by the number of identified buyer links for retailer r , not by the count in each possible $r \times h$ cell. [Table B.2](#) shows that these effective sample sizes are large: even in the lowest percentile bin, the average number of identified buyer IDs is 139, rising to 189 at the 25th percentile, 983 at the median, 2,010 at the 75th percentile, and 3,290 at the 90th percentile. Thus, the catchment objects are estimated from hundreds or thousands of observed customer links, not from a handful of realized links.

Formally, let D_i indicate whether invoice i includes a buyer ID, let H_i denote the buyer’s residence neighborhood, and let v_i denote the transaction value. The sales-weighted empirical catchment weight is

$$\hat{w}_{rh} = \frac{\sum_{i \in \mathcal{I}_r} D_i v_i \mathbf{1}\{H_i = h\}}{\sum_{i \in \mathcal{I}_r} D_i v_i},$$

with the analogous expression for customer-count weights obtained by setting $v_i = 1$. If ID disclosure is mean independent of residence within retailer, then $\hat{w}_{rh}$ is a consistent estimate of w_{rh} . More generally, for any bounded neighborhood-level object x_h ,

$$\hat{x}_r \equiv \sum_h \hat{w}_{rh} x_h$$

has sampling variance of order $1/n_r$, where n_r is the number of identified buyer links for retailer r . With sales weights, the same logic applies with n_r replaced by the

effective sample size

$$n_r^{\text{eff}} \equiv \frac{(\sum_{i \in \mathcal{I}_r} D_i v_i)^2}{\sum_{i \in \mathcal{I}_r} D_i v_i^2}.$$

This is the sense in which the low-dimensional catchment average is far less granular than a sparse origin–destination matrix.

The remaining concern is non-classical missingness: ID disclosure could be systematically related to the same retailer characteristics that drive our analysis. Figure B.1 directly assesses this possibility by plotting the share of transactions with customer IDs against retailer market shares, catchment-area income, and catchment-area exposure. The coverage share is flat across these observables. Hence, identified invoices are not concentrated among retailers with particular levels of market power, income, or exposure. To rationalize our results through selective non-disclosure, missing IDs would have to vary within retailer in a product-specific way that is correlated with pre-policy foreign shares and the HS-4 tariff changes, and do so more strongly for high-market-share retailers, while leaving no trace in the coverage diagnostics in Figure B.1. This is a restrictive alternative. Under the more natural case of residual sampling error in w_{rh} , the resulting measurement error in

$$\widehat{IC}_{rg} = IC_{rg} + \nu_{rg}$$

would attenuate the estimated exposure effects and blur the relationship between price responses and market shares, but it would not generate the sharp upper-tail price responses that we document.

B.3 Alternative Catchment-Area Constructions

This appendix provides two complementary validation exercises for the catchment-area construction. Both are designed to address the concern that customer IDs are observed only for a subset of retail invoices. We compare the baseline weights w_{rh} , based on customers’ residence neighborhoods, to two alternative constructions. In each case, we reconstruct the catchment weights, recompute retailer–product exposure IC_{rg} , recompute catchment-area market shares, and re-estimate both the reduced-form and markup-elasticity exercises.

PPML gravity reconstruction. The first alternative predicts retailer–neighborhood weights using a PPML gravity model estimated on high-ID-coverage retailers. Namely, we consider retailers with a share of transactions with customer identifiers in the 80th percentile of the distribution or above. Let $\mathcal{R}^{80}$ denote this set of retailers with high ID coverage. For these retailers, the identified-invoice data provide a close approximation to the geography of the retailer’s full customer base. Let y_{rh}^{ID} denote identified sales from retailer r to customers residing in neighborhood h . We estimate

$$\mathbb{E}[y_{rh}^{ID} \mid X_{rh}] = \exp(\alpha_r + \delta_h + X'_{rh}\beta), \quad r \in \mathcal{R}^{80},$$

where α_r is a retailer fixed effect, δ_h is a neighborhood fixed effect, and X_{rh} contains pre-policy observables. The predictor set includes bilateral distance between the retailer and the neighborhood, neighborhood income, population, and formal-employment share, as well as retailer characteristics such as total sales, average prices, and the premium-product share. Retailer characteristics enter through interactions with distance and neighborhood characteristics, allowing larger retailers to draw from farther neighborhoods and premium retailers to draw more heavily from higher-income neighborhoods. The retailer fixed effects absorb the scale of each training retailer and are not used for prediction.

The predicted weights for any retailer r are obtained by normalizing predicted flows across neighborhoods:

$$w_{rh}^{PPML} = \frac{\exp(\widehat{\delta}_h + X'_{rh}\widehat{\beta})}{\sum_{k \in \mathcal{H}} \exp(\widehat{\delta}_k + X'_{rk}\widehat{\beta})}.$$

Thus, the procedure predicts the weights directly, preserves the weighted-average structure of the baseline catchment construction, and avoids choosing a threshold for whether a neighborhood belongs to a catchment area.

Employee-residence reconstruction. The second alternative relies on employer–employee records. For each retailer r , we observe the residence neighborhoods of its employees. Let Emp_{rh} denote the number of employees of retailer r residing in neighborhood h . We define

$$w_{rh}^{Emp} = \frac{Emp_{rh}}{\sum_{k \in \mathcal{H}} Emp_{rk}}.$$

This construction uses employees’ residences as a proxy for the geography of the retailer’s customer base. The intuition is that employees provide information about the spatial footprint of the retailer: larger retailers have more employees and therefore a broader empirical support, while multi-establishment retailers tend to have employees spread across multiple parts of the country. This proxy is useful because employer–employee records are often available in administrative data even when customer locations are not.

Agreement with baseline weights. Table B.3 reports the correlation between the alternative weights and the baseline customer-residence weights. The PPML reconstruction is strongly correlated with the baseline, with a correlation of 0.77. The employee-residence construction is also highly correlated, with a correlation of 0.71. Thus, although the two alternatives use very different sources of information, both recover a similar geography of retailer catchment areas.

Table B.3: Correlation with Baseline Catchment Weights

Alternative catchment construction	Corr. with baseline w_{rh}
PPML gravity reconstruction, w_{rh}^{PPML}	0.77
Employee residences, w_{rh}^{Emp}	0.71

Notes: The table reports correlations between the baseline catchment weights, constructed from customers' residence neighborhoods in retail invoices, and the two alternative catchment-weight constructions. The PPML reconstruction predicts retailer-neighborhood weights using distance, retailer characteristics, neighborhood characteristics, and their interactions. The employee-residence construction uses the residence neighborhoods of each retailer's employees.

Recomputing exposure and market shares. For each alternative construction $m \in \{PPML, Emp\}$, we recompute the foreign share in retailer r 's catchment area for product category g :

$$\bar{S}_{fg0}^{r,m} = \sum_{h \in \mathcal{H}} w_{rh}^m S_{h,fg0},$$

and define the corresponding exposure measure $IC_{rg}^m = \bar{S}_{fg0}^{r,m} \Delta \tau_g$. We also recompute catchment-area market shares using the same weights: $S_{rg0}^m = \sum_{h \in \mathcal{H}} w_{rh}^m S_{rh}^h$. For retailer-level labor outcomes, we aggregate exposure across products using pre-policy sales weights, obtaining IC_r^m .

B.4 Results with Alternative Catchment-Area Constructions

Reduced-form robustness. Tables B.4 and B.5 report the two-year reduced-form results using the alternative catchment-area constructions. The qualitative pattern is the same as in the baseline. More exposed retailer-HS-4 cells increase prices and sales, while input-cost and labor-cost outcomes do not exhibit corresponding increases. For the main price and sales outcomes, the baseline estimates lie between the two alternative constructions.

Table B.4: Reduced-Form Impacts Using PPML-Reconstructed Catchments

	Retailer-HS-4 outcomes		Input costs		Labor outcomes	
	$\Delta \ln(\text{Price}_{rg})$ (1)	$\Delta \ln(\text{Sales}_{rg})$ (2)	$\Delta \ln(\text{Import Price}_{rg})$ (3)	$\Delta \ln(\text{Purchase Price}_{rg})$ (4)	$\Delta \ln(\text{Employment}_r)$ (5)	$\Delta \ln(\text{Wage}_r)$ (6)
IC_{rg}^{PPML}	0.053*** (0.006)	0.023*** (0.004)	0.002 (0.001)	0.005 (0.004)		
IC_r^{PPML}					-0.005 (0.006)	-0.002 (0.004)
Retailer FE	✓	✓	✓	✓		
HS-4 FE	✓	✓	✓	✓		
$\bar{S}_{fg0}^{r,PPML}$	✓	✓	✓	✓		
Retailer controls $\mathbf{X}_{rT}^{PPML}$					✓	✓
Main HS-4 FE					✓	✓
Industry FE					✓	✓
Observations	320,753	320,753	7,764	162,793	20,752	20,752

Notes: The table reports two-year long-difference regressions using catchment weights reconstructed from the PPML gravity model. Columns (1)–(2) are retailer-HS-4 price and sales regressions. Columns (3)–(4) are input-cost regressions; column (3) uses the unit import price $\log(\text{dollar}_M/\text{quantity}_M)$, while column (4) uses the purchase-price outcome from firm-to-firm input purchases. Columns (1)–(4) include retailer and HS-4 fixed effects and control for $\bar{S}_{fg0}^{r,PPML}$. Columns (5)–(6) are retailer-level labor regressions using the sales-weighted exposure IC_r^{PPML} . The vector of retailer controls is $\mathbf{X}_{rT}^{PPML} = (\bar{S}_{fg0}^{r,PPML}, \bar{\Delta\tau}_r, \Delta_T \ln R_r)'$, where $\bar{S}_{fg0}^{r,PPML}$ is the retailer-level baseline foreign-share control, $\bar{\Delta\tau}_r$ is the sales-weighted tariff-change component, and $\Delta_T \ln R_r$ is the change in total retailer revenues. Standard errors, clustered by HS-4, are in parentheses; except for columns (5)–(6), which cluster by retailer. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.5: Reduced-Form Impacts Using Employee-Residence Catchments


	Retailer-HS-4 outcomes		Input costs		Labor outcomes	
	$\Delta \ln(\text{Price}_{rg})$ (1)	$\Delta \ln(\text{Sales}_{rg})$ (2)	$\Delta \ln(\text{Import Price}_{rg})$ (3)	$\Delta \ln(\text{Purchase Price}_{rg})$ (4)	$\Delta \ln(\text{Employment}_r)$ (5)	$\Delta \ln(\text{Wage}_r)$ (6)
IC_{rg}^{Emp}	0.009*** (0.001)	0.007*** (0.001)	-0.004 (0.008)	0.001 (0.002)		
IC_r^{Emp}					-0.006** (0.003)	0.003 (0.002)
Retailer FE	✓	✓	✓	✓		
HS-4 FE	✓	✓	✓	✓		
$\bar{S}_{fg0}^{r,Emp}$	✓	✓	✓	✓		
Retailer controls $\mathbf{X}_{rT}^{Emp}$					✓	✓
Main HS-4 FE					✓	✓
Industry FE					✓	✓
Observations	262,674	262,674	7,586	136,525	20,752	20,752

Notes: The table reports two-year long-difference regressions using catchment weights constructed from employees' residence neighborhoods. Columns (1)–(2) are retailer-HS-4 price and sales regressions. Columns (3)–(4) are input-cost regressions; column (3) uses the unit import price $\log(\text{dollar}_M/\text{quantity}_M)$, while column (4) uses the purchase-price outcome from firm-to-firm input purchases. Columns (1)–(4) include retailer and HS-4 fixed effects and control for $\bar{S}_{fg0}^{r,Emp}$. Columns (5)–(6) are retailer-level labor regressions using the sales-weighted exposure IC_r^{Emp} . The vector of retailer controls is $\mathbf{X}_{rT}^{Emp} = (\bar{S}_{fg0}^{r,Emp}, \bar{\Delta\tau}_r, \Delta_T \ln R_r)'$, where $\bar{S}_{fg0}^{r,Emp}$ is the retailer-level baseline foreign-share control, $\bar{\Delta\tau}_r$ is the sales-weighted tariff-change component, and $\Delta_T \ln R_r$ is the change in total retailer revenues. Standard errors, clustered by HS-4, are in parentheses; except for columns (5)–(6), which cluster by retailer. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Markup-elasticity robustness. Finally, we ask whether the estimated heterogeneity in markup elasticities is sensitive to the catchment-area construction. For each set of weights, we reconstruct exposure and market shares, estimate the same IV price-share elasticities used in Figure 8, and plot the resulting markup-elasticity schedule across the market-share distribution.

Figure B.2 shows that the three schedules are extremely close. All three constructions imply modest markup elasticities at low and intermediate market shares and a sharp increase in the upper tail. This is the key robustness result: the non-linear markup-elasticity pattern is not an artifact of observing customer IDs in retail invoices. It is a salient feature of the data that remains when catchment areas are reconstructed from a gravity model or proxied using employee residence neighborhoods.

Figure B.2: Markup Elasticities Across Catchment-Area Constructions



Notes: The figure plots the IV-estimated markup-elasticity schedule across the market-share distribution using three catchment-area constructions: the baseline customer-residence weights from retail invoices, the PPML gravity reconstruction, and the employee-residence proxy. For each construction, we recompute exposure and catchment-area market shares before estimating the price-share elasticity by market-share bin.

Overall, both validation exercises support the baseline construction. The alternative weights are strongly correlated with the customer-residence weights, the reduced-form price and sales responses remain positive and statistically significant, input-cost responses remain close to zero, and the heterogeneous markup-elasticity schedule is nearly unchanged across methods. The employee-residence exercise is also useful beyond our setting, since employer-employee records are often available even when customer locations are not.

C Reduced Form: Additional Robustness

C.1 Alternative Horizon

Table C.1 reports reduced-form specifications as in Tables 1-3, but using a one-year post-reform horizon. The documented patterns align across horizons.

Table C.1: Reduced-Form Impacts: One-Year Horizon

	Retailer-HS-4 outcomes		Retailer outcomes		Input costs		Substitution
	$\Delta \ln(\text{Price}_{rg})$ (1)	$\Delta \ln(\text{Sales}_{rg})$ (2)	$\Delta \ln(\text{Employment}_r)$ (3)	$\Delta \ln(\text{Wage}_r)$ (4)	$\Delta \ln(\text{Import Price}_{rg})$ (5)	$\Delta \ln(\text{Purchase Price}_{rg})$ (6)	Δs_{hg} (7)
IC_{rg}	0.018*** (0.004)	0.013*** (0.003)			0.007 (0.018)	0.001 (0.003)	
IC_r			0.005 (0.004)	0.005 (0.003)			
IC_{hg}							-0.00018 (0.00023)
Retailer FE	✓	✓			✓	✓	
HS-4 FE	✓	✓			✓	✓	✓
$\bar{S}_{fg0}^r$	✓	✓			✓	✓	
Retailer controls $\mathbf{X}_{rT}$			✓	✓			
Main HS-4 FE			✓	✓			
Industry FE			✓	✓			
$\bar{S}_{fg0}^h$ FE							✓
Neighborhood FE							✓
Observations	342,573	342,573	21,171	21,171	8,161	173,591	595

Notes: Columns (1)–(2) are retailer–HS-4 regressions and include retailer and HS-4 fixed effects. Columns (3)–(4) are retailer-level regressions using the sales-weighted exposure IC_r . The vector of retailer controls is $\mathbf{X}_{rT} = (\bar{S}_{fg0}^r, \bar{\Delta\tau}_r, \Delta_T \ln R_r)'$, where $\bar{S}_{fg0}^r$ is the retailer-level baseline foreign-share control, $\bar{\Delta\tau}_r$ is the sales-weighted tariff-change component, and $\Delta_T \ln R_r$ is the change in total retailer revenues. Columns (5)–(6) are retailer–HS-4 input-cost regressions. Column (5) uses the unit import price $\log(\text{dollar}_M/\text{quantity}_M)$, while column (6) uses the purchase-price outcome from firm-to-firm input purchases. Column (7) considers the change in the expenditure share of category g on the neighborhood–HS-4 exposure measure IC_{hg} . Standard errors, clustered by HS-4, are in parentheses; except for columns (3)–(4), which cluster by retailer. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C.2 Results by HS-6 Category

Tables C.2 and C.3 repeat the main reduced form regressions using HS-6 instead of HS-4 product categories. The documented patterns align across product categories.

Table C.2: Reduced-Form Using HS-6 Categories: Baseline Two-Year Horizon

	Retailer-HS-6 outcomes		Input costs	
	$\Delta \ln(\text{Price}_{rg})$	$\Delta \ln(\text{Sales}_{rg})$	$\Delta \ln(\text{Import Price}_{rg})$	$\Delta \ln(\text{Purchase Price}_{rg})$
	(1)	(2)	(3)	(4)
IC_{rg}^{HS6}	0.013*** (0.001)	0.006* (0.003)	-0.042* (0.023)	-0.001 (0.004)
Retailer FE	✓	✓	✓	✓
HS-6 FE	✓	✓	✓	✓
$\bar{S}_{fg0}^r$	✓	✓	✓	✓
Observations	425,731	425,731	9,656	206,291

Notes: The table reports two-year long-difference regressions at the retailer-HS-6 level. Columns (1)–(2) are price and sales regressions. Columns (3)–(4) are input-cost regressions; column (3) uses the unit import price $\log(\text{dollar}_M/\text{quantity}_M)$, while column (4) uses the purchase-price outcome from firm-to-firm input purchases. All regressions include retailer fixed effects, HS-6 fixed effects, and the baseline foreign-share control $\bar{S}_{fg0}^r$, where g denotes an HS-6 product category. Standard errors, clustered by HS-6, are in parentheses. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.3: Reduced-Form Using HS-6 Product Categories: One-Year Horizon


	Retailer-HS-6 outcomes		Input costs	
	$\Delta \ln(\text{Price}_{rg})$	$\Delta \ln(\text{Sales}_{rg})$	$\Delta \ln(\text{Import Price}_{rg})$	$\Delta \ln(\text{Purchase Price}_{rg})$
	(1)	(2)	(3)	(4)
IC_{rg}^{HS6}	0.012** (0.005)	0.005** (0.003)	-0.014 (0.027)	0.007 (0.004)
Retailer FE	✓	✓	✓	✓
HS-6 FE	✓	✓	✓	✓
$\bar{S}_{fg0}^r$	✓	✓	✓	✓
Observations	482,244	482,244	10,302	219,669

Notes: The table reports one-year long-difference regressions at the retailer-HS-6 level. Columns (1)–(2) are price and sales regressions. Columns (3)–(4) are input-cost regressions; column (3) uses the unit import price $\log(\text{dollar}_M/\text{quantity}_M)$, while column (4) uses the purchase-price outcome from firm-to-firm input purchases. All regressions include retailer fixed effects, HS-6 fixed effects, and the baseline foreign-share control $\bar{S}_{fg0}^r$, where g denotes an HS-6 product category. Standard errors, clustered by HS-6, are in parentheses. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D National Market Shares by HS-4 and HS-6

D.1 National Shares by HS-4

Figure D.1: Markup Elasticities Based on National Shares by HS-4 Category



Notes: The figure repeats the IV price–share elasticity exercise after replacing catchment-area market shares with national retailer–HS-4 shares. A national share is defined as retailer r 's baseline domestic sales in HS-4 category g divided by total national domestic sales in that category. The resulting schedule is estimated mostly from very low market shares, because national shares collapse geographically segmented pricing markets into a single national product market. As shown in Table D.1, the national-share distribution contains essentially no mass in the high-share region where the theory predicts markup elasticities should be largest.


Table D.1: Distribution of National Retailer–HS-4 Market Shares

	Percentile						
	10th	25th	50th	75th	90th	95th	99th
National share (%)	0.00019	0.00068	0.00318	0.01847	0.104	0.328	3.306

Notes: The table reports the distribution across retailer–HS-4 cells of each retailer's national market share within an HS-4 category. Shares are expressed in percent. The sample contains 864,053 retailer–HS-4 cells. The mean national market share is 0.196%, and the standard deviation is 2.133%. Only 48 retailer–HS-4 cells have a national market share above 90%.

D.2 National Shares by HS-6

Figure D.2: Markup Elasticities Based on National Shares by HS-6 Category



Notes: The figure repeats the IV price–share elasticity exercise after replacing catchment-area market shares with national retailer–HS-6 shares. A national share is defined as retailer r ’s baseline domestic sales in HS-6 category g divided by total national domestic sales in that category. Just as with HS-4 categories, the resulting schedule is estimated mostly from very low market shares, because national shares collapse geographically segmented pricing markets into a single national product market. As shown in Table D.2, the national-share distribution contains essentially no mass in the high-share region where the theory predicts markup elasticities should be largest.

Table D.2: Distribution of National Retailer–HS-6 Market Shares

	Percentile						
	10th	25th	50th	75th	90th	95th	99th
National share (%)	0.00035	0.00136	0.00666	0.0392	0.225	0.739	7.361

Notes: The table reports the distribution across retailer–HS-6 cells of each retailer’s national market share within an HS-6 category. Shares are expressed in percent. The sample contains 1,106,379 retailer–HS-6 cells. The mean national market share is 0.373%, and the standard deviation is 3.210%.


E Results for Single-Establishment Retailers

Table E.1: Reduced-Form Impacts for Single-Establishment Retailers

	Retailer–HS-4 outcomes		Input costs	
	$\Delta \ln(\text{Price}_{rg})$	$\Delta \ln(\text{Sales}_{rg})$	$\Delta \ln(\text{Import Price}_{rg})$	$\Delta \ln(\text{Purchase Price}_{rg})$
	(1)	(2)	(3)	(4)
IC_{rg}	0.028*** (0.006)	0.018*** (0.004)	−0.030 (0.022)	0.002 (0.004)
Retailer FE	✓	✓	✓	✓
HS-4 FE	✓	✓	✓	✓
$\bar{S}_{fg0}^r$	✓	✓	✓	✓
Single-establishment sample	✓	✓	✓	✓
Observations	292,146	292,146	5,101	141,819

Notes: The table reports two-year long-difference regressions for the subsample of single-establishment retailers. Columns (1)–(2) are retailer–HS-4 price and sales regressions. Columns (3)–(4) are input-cost regressions; column (3) uses the unit import price $\log(\text{dollar}_M/\text{quantity}_M)$, while column (4) uses the purchase-price outcome from firm-to-firm input purchases. All regressions include retailer fixed effects, HS-4 fixed effects, and the baseline foreign-share control $\bar{S}_{fg0}^r$. Standard errors, clustered by HS-4, are in parentheses. Significance is denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure E.1: Markup Elasticities for Single-Establishment Retailers



Notes: The figure repeats the IV price–share elasticity exercise for the subsample of single-establishment retailers. The empirical series reports the estimated elasticity of prices with respect to market shares across bins of the catchment-area market-share distribution. Exposure IC_{rg} is used as an instrument for changes in retailer–HS-4 market shares, and the regressions include the same controls and fixed effects as in the baseline specification. The shaded region denotes the confidence interval.

F Model Details

F.1 Derivation of the Pricing Rule

The first-order condition is

$$0 = Y_{rgt} + (P_{rgt} - A_{rgt}) \sum_{h \in Z_r} \frac{\partial Y_{rgt}^h}{\partial P_{rgt}}. \quad (22)$$

Under the nested CES demand system in [equation \(8\)](#)–[equation \(9\)](#), the demand elasticity faced by retailer r in neighborhood h is

$$-\frac{\partial \log Y_{rgt}^h}{\partial \log P_{rgt}} = \sigma(1 - S_{hrgt}) + \gamma S_{hrgt},$$

where S_{hrgt} is retailer r 's expenditure share in neighborhood h and product group g . Therefore,

$$\frac{\partial Y_{rgt}^h}{\partial P_{rgt}} = -\frac{Y_{rgt}^h}{P_{rgt}} [\sigma(1 - S_{hrgt}) + \gamma S_{hrgt}].$$

Substituting this expression into [equation \(22\)](#) gives

$$1 = \frac{P_{rgt} - A_{rgt}}{P_{rgt}} \sum_{h \in Z_r} \frac{Y_{rgt}^h}{Y_{rgt}} [\sigma(1 - S_{hrgt}) + \gamma S_{hrgt}].$$

Because the retailer sets a uniform price across neighborhoods,

$$\frac{Y_{rgt}^h}{Y_{rgt}} = \frac{E_{hrgt}}{\sum_{\ell \in Z_r} E_{\ell rgt}},$$

where $E_{hrgt} = P_{rgt} Y_{rgt}^h$. This motivates the catchment-area market share

$$S_{rgt} \equiv \sum_{h \in Z_r} \omega_{hrgt} S_{hrgt}, \quad \omega_{hrgt} \equiv \frac{E_{hrgt}}{\sum_{\ell \in Z_r} E_{\ell rgt}}. \quad (23)$$


Using this definition, the first-order condition becomes

$$1 = \frac{P_{rgt} - A_{rgt}}{P_{rgt}} [\sigma(1 - S_{rgt}) + \gamma S_{rgt}],$$

which yields the pricing rule in [equation \(11\)](#).


G Model-Relevant Empirical Regularities

Figure G.1: Domestic Sales Across the Market-Share Distribution




Notes: The figure reports the share of domestic retail sales accounted for by retailers in each bin of the market-share distribution. Market shares are measured at the retailer–HS-4 category level using the baseline catchment-area market definition.

Figure G.2: Input Prices Across the Market-Share Distribution



Notes: The figure repeats the reduced-form exercise for input prices across bins of the baseline catchment-area market-share distribution. The outcome is a retailer–HS-4 input-price index constructed as a purchase-weighted average of the inputs bought by retailer r in category g , combining both domestic firm-to-firm purchases and direct imports. This pooled measure provides coverage across the market-share distribution, since retailers typically source a given retailer–HS-4 cell either domestically or through imports. The regressions include retailer fixed effects, HS-4 fixed effects, and the baseline foreign-share control S_{fg0}^r . The shaded region denotes confidence intervals. Input-price responses are close to zero over most of the distribution and do not mirror the steep upper-tail pattern in retail prices.

Figure G.3: Distributional Effects: The Cost of Remoteness



Notes: The figures show outcomes against distance to the nearest international port or airport. Panel (a) shows that the largest exposure occurs in most connected areas. Panel (b) shows that the highest market shares occur in most remote areas.