Sourcing Networks, Foreign Shocks, and Imported Inflation in an Emerging Market

Monica Morlacco*

September 8, 2026

Abstract

Terms-of-trade fluctuations are often modeled as primitive external shocks in small open economies. Using Colombian import transactions from 2001–2025, I show that import-price inflation also reflects the granular structure of sourcing. A *granular residual* capturing the influence of large relationships accounts for 70 percent of dollar import-price fluctuations, primarily because these relationships concentrate expenditure on products experiencing large common price movements. Large relationships also reprice differently within products, reflecting supplier costs rather than markup adjustment, and the pass-through of foreign cost shocks declines with relationship size. These microeconomic origins matter little for the contemporaneous transmission of imported inflation to domestic producer prices but predict persistence: among continuing links, common product-price increases reverse roughly twice as much as relationship-specific increases. More broadly, the same foreign shocks can generate different national import-price inflation under different sourcing configurations.

Keywords: Imported inflation; terms of trade; granular fluctuations; firm-to-firm trade; global shocks; pass-through.
JEL codes: E31, F14, F41, L11.

*I thank Andrei Levchenko for helpful feedback. This research received support only from the University of Southern California. The author has no financial relationships or other potential conflicts of interest relevant to this paper. All errors are my own. Email: morlacco@usc.edu.

1 Introduction

Fluctuations in the terms of trade are a central policy concern in emerging markets because they shape inflation, external adjustment, and economic activity. Workhorse models of small open economies typically treat these fluctuations as an exogenous process (Mendoza, 1995; Schmitt-Grohé and Uribe, 2018). In the data, however, the terms of trade are not primitive disturbances. They are constructed from import and export price indexes that aggregate price changes across individual buyer–supplier relationships using trade-value weights. These relationships form highly concentrated networks, in which a small number of links account for a large share of trade (Eaton et al., 2011; Bernard et al., 2018). When trade is granular, shocks to dominant links can shape aggregate fluctuations and trade patterns (Gabaix, 2011; Gaubert and Itskhoki, 2021). Applied to prices, this insight implies that even when the underlying foreign shocks are exogenous, their mapping into observed import and export prices may depend on idiosyncratic patterns of exposure and transmission.

This paper asks two related questions. First, what are the microeconomic origins of import-price fluctuations? Second, does their microeconomic composition shape domestic transmission and persistence? I address these questions using transaction-level data on Colombian imports from 2001 to 2025, a period spanning the global trade collapse, a sharp depreciation of the peso, and the post-pandemic inflation surge. I show that Colombia's concentrated sourcing network materially affects realized import-price inflation through two channels. An exposure channel determines the weights placed on shocks to individual products, while a repricing channel captures the effect of relationship-specific price changes on the aggregate index. Although both channels are quantitatively important, exposure accounts for the larger share in this emerging market setting. The differential repricing behavior of large relationships primarily reflects differences in costs rather than markups, together with heterogeneous pass-through of cost shocks.

This composition matters for the macroeconomic interpretation of import-price movements. Common product-price movements and relationship-specific repricing pass through similarly to domestic producer prices within a year, making the aggregate index a useful measure of current imported cost pressure. Their subsequent paths differ, however: common product-price increases reverse roughly twice as much as relationship-specific increases. Measured terms-of-trade fluctuations therefore reflect both the external disturbance and the country-specific mapping of that disturbance through expenditure weights and bilateral pass-through. Treating them as primitive shocks conflates the disturbance with its transmission. More generally, counterfactuals that change where and from whom a country imports must allow the import-price process to change with the sourcing network.

The data comprise the universe of Colombian import declarations from 2001 to 2025, observed at a monthly frequency. I map each Colombian importer to its foreign exporter and the variety imported, defined as a ten-digit Harmonized System (HS10) code paired with the production origin. I construct a relationship-level import-price index by aggregating price changes within continuing importer–exporter–variety relationships and weighting them using lagged import expenditures. Holding relationships fixed over time separates repricing from changes in the extensive margin of sourcing and makes it possible to trace how concentrated expenditure maps transaction-level price changes into aggregate inflation. Despite substantial turnover in these relationships, the resulting index closely tracks both conventional product–country indexes constructed from the same customs data and Colombia's official import producer price index. Preserving relationship-level detail thus captures the bulk of aggregate import-price fluctuations while allowing their microeconomic origins to be identified.

Colombian imports are highly concentrated. The largest 1 percent of continuing relationships account for 54 percent of import value, and the largest decile accounts for nearly 90 percent. Despite tens of thousands of active relationships, the effective number is only a few hundred. To quantify the aggregate importance of this concentration, I construct the *granular residual*, defined as the difference between value-weighted import-price inflation and the equal-weighted mean price change across relationships. This measure follows the approach pioneered by Gabaix (2011) and recently applied to retail inflation by Alvarez-Blaser et al. (2025). A covariance decomposition attributes 70 percent of the fluctuations in four-quarter, dollar-denominated import-price inflation to the granular residual.

Most of this contribution operates through *granular product exposure*: large relationships are disproportionately concentrated in products experiencing unusually large common price movements. *Granular repricing* arises when prices in large relationships adjust differently from those in other relationships involving the same product. This channel accounts for a smaller but economically meaningful share of aggregate fluctuations, particularly in energy and differentiated goods. Overall, product exposure contributes 54 of the 70 percentage points attributed to the granular residual, while granular repricing contributes the remaining 16 points. This finding resonates with Alvarez-Blaser et al. (2025), who find that common forces account for a larger share of retail inflation in emerging markets than in advanced economies. The evidence here shows that common shocks can nevertheless have granular aggregate effects: when sourcing is concentrated, unequal expenditure weights transform a common vector of product-price movements into a country-specific import-price index.

Concentration is neither inherently inflationary nor stabilizing. Its aggregate effect depends on the covariance between expenditure weights and price changes. The two largest

inflationary episodes in the sample illustrate this point. During the 2014–16 depreciation, the collapse in oil and other commodity prices weakened the peso but simultaneously lowered the dollar prices of many imports. Large relationships were disproportionately concentrated in products experiencing the steepest price declines and, within those products, obtained larger price reductions than the average importer. Relative to an equal-weighted benchmark, concentration reduced peso import-price inflation by 7.6 percentage points. During the 2021–22 global inflation surge, these same margins operated in reverse: dominant relationships were concentrated in products with rapidly rising prices, increasing peso import-price inflation by 5.4 percentage points.

This contrast raises a natural question: did the effect of concentration change because the sourcing network was reconfigured, or because different shocks hit a persistent set of expenditure weights? I distinguish these mechanisms using fixed-share counterfactuals that apply the same observed product–origin price changes to contemporaneous and pre-episode expenditure weights. Despite substantial relationship turnover, changes in product–origin exposure had modest effects in both episodes. Sourcing reallocation reduced import-price inflation by 0.43 percentage points in 2014–16 and by 0.32 percentage points in 2021–22. Despite extensive link turnover, Colombia's dominant sourcing exposure remained stable through two major global shocks.[1] In both episodes, relationship churn did not act as a source of aggregate vulnerability to foreign shocks.

With aggregate exposure largely stable, I next ask why large relationships reprice differently within products. A natural explanation is endogenous markup adjustment. Variable markups are central to theories of international pass-through (Atkeson and Burstein, 2008; Amiti et al., 2019; Gopinath et al., 2020) and feature prominently in the debate over post-pandemic inflation, even though evidence on their role in shaping aggregate price dynamics is mixed (Weber and Wasner, 2023; Conlon, 2026). Observed price changes, however, combine movements in marginal costs and markups. A common challenge of price-based decompositions alone is therefore the inability to separate whether granular repricing reflects supplier-side cost shocks or adjustments in bilateral margins (Amiti et al., 2024; Alvarez-Blaser et al., 2025).

The relationship-level data allow me to separate these forces. To do so, I use the bilateral bargaining model of Alviarez et al. (2025). The model maps observed purchase and sales shares into relationship-specific markups and attributes the remaining price variation to marginal costs. I find that model-implied markup levels vary substantially across relationships but change little over time because the market shares governing them are highly persistent.

[1]This pattern echoes Bricongne et al. (2012), who show that smaller exporters responded to the global trade collapse through entry and exit, while large firms preserved their links and adjusted primarily along the intensive margin.

Markup changes account for only 2.1% of the covariance with relationship-level price changes and virtually none of aggregate import-price inflation. This pattern is consistent with the bottom-up markup mechanism in Burstein et al. (2025): markups respond when shocks reallocate market shares, whereas common shocks can move costs and prices without materially changing relative shares. The evidence therefore points to supplier-level cost movements, rather than markup adjustment, as the main source of granular repricing. Large relationships matter primarily because they concentrate the economy's exposure to cost shocks, not because they adjust markups more aggressively.

The evidence that marginal costs drive granular repricing motivates a direct test using an observable foreign cost shifter. Following the exchange-rate pass-through literature (Berman et al., 2012; Amiti et al., 2014), I estimate the pass-through of movements in suppliers' currencies relative to the dollar into dollar-denominated transaction prices. Conditional on product–quarter fixed effects, pass-through is broadly constant across the lower half of the relationship-size distribution but declines sharply through the upper half. A 10% appreciation of the supplier's currency raises unit values by approximately 2.4% in the smallest decile, compared with only 0.8% in the largest. This result identifies a second role for the sourcing network. In addition to determining which foreign shocks receive the greatest weight, the network shapes how strongly those shocks pass through into border prices.

The macroeconomic relevance of these microeconomic distinctions depends on the horizon. In the short run, composition matters little for domestic transmission. Using Colombia's input-use structure to map common product-price movements and relationship-specific repricing into sectoral input costs, I find that a one-percentage-point increase in imported-input costs is associated with approximately 0.56 percentage points of producer-price inflation within a year. The estimate is virtually identical across the two components, and I find no differential response of production, sales, or employment. At longer horizons, however, composition predicts persistence. Tracking the same buyer–supplier links forward, I find that common product-price increases are followed by reversals of 16–21%, compared with 6–11% for relationship-specific increases. The specification conditions on the initial import-price increase, so the estimates compare movements of similar initial magnitude and isolate differences in their subsequent dynamics. The gap persists among the largest 1% of relationships and remains visible after aggregation to product groups.

Together, these results show that the terms of trade cannot be separated from the structure of trade itself. Import prices depend not only on foreign price changes, but also on what the country imports, from whom, and how prices adjust within each trading relationship. Empirical work that treats terms-of-trade movements as exogenous shocks must therefore

ensure that sourcing weights and pricing responses are not themselves driven by domestic conditions.

The paper contributes to the literature on granular fluctuations. Previous work shows that shocks to large firms and firm–destination links survive aggregation and shape trade flows and macroeconomic outcomes (Gabaix, 2011; Di Giovanni et al., 2014; Freund and Pierola, 2015; Gaubert and Itskhoki, 2021; di Giovanni et al., 2022), while Alvarez-Blaser et al. (2025) establish the granular origins of retail inflation. I extend this logic to border prices and to the importer–exporter–variety relationship, a unit of observation closely tied to bilateral price formation. This approach separates the exposure of dominant links to common product-price movements from their differential repricing within products. It also provides evidence that differential repricing primarily reflects supplier-side cost movements rather than markup adjustment, a distinction that price-only decompositions cannot make.

The paper also contributes to research on cross-border shock propagation and terms-of-trade shocks. Auer et al. (2019) show how input-output linkages propagate cost shocks across countries, while Amiti et al. (2024) decompose imported inflation into global, exporter-specific, and importer-specific components across country–product trade flows. I show that the import-price movement entering the domestic production network depends on granular sourcing. Whereas Di Pace et al. (2025) distinguish between movements originating in import and export prices, I show that economically relevant heterogeneity remains within import prices themselves through their granular composition. These findings challenge the common practice of treating terms-of-trade movements as exogenous shocks.

The remainder of the paper proceeds as follows. Section 2 describes the data and constructs the relationship-level price index. Section 3 documents concentration, provides a granular accounting of import-price inflation, and studies the contribution of extensive margin changes in sourcing. Section 4 investigates differential repricing and the heterogeneous pass-through of supplier-currency shocks. Section 5 examines the macroeconomic consequences of the granular composition of import-price inflation. Section 6 concludes.

2 Data and Measurement

This section describes the monthly customs panel and the construction of import price indexes. I define a relationship as a Colombian importer, a foreign exporter, and a variety, an HS10 product code paired with its country of production origin. A relationship is *continuing* at horizon k if the same importer–exporter–variety cell is observed in both $m-k$ and m. I denote relationships by r throughout.

2.1 Data Sources and Panel Construction

The main data source is the universe of Colombian merchandise import declarations recorded by the customs authority, the *Dirección de Impuestos y Aduanas Nacionales* (DIAN). I reconstruct these declarations from DIAN's monthly statistical archives covering 2001–2025. Each declaration includes the Colombian importer's tax identifier (*Número de Identificación Tributaria*, or NIT), the foreign exporter's name and reported city and country, a ten-digit Harmonized System (HS10) product code, the country of production origin, FOB value in U.S. dollars, and quantity in kilograms.

Colombian importers have unique identifiers and can therefore be tracked over time. Foreign exporters, by contrast, lack unique identifiers. Exporter names are recorded as free text, truncated at forty characters, and often vary across declarations because of differences in spelling, abbreviations, and formatting. I construct longitudinal exporter identifiers by standardizing reported names and linking likely matches across periods. Appendix C.1 describes the matching procedure and validation exercises.

I also harmonize product codes across revisions of the HS classification, which occur every five years, using official correspondence tables following Pierce and Schott (2009). Countries of production and exporter origin are standardized using ISO country codes.

I restrict the sample to observations with positive value and quantity and valid importer, exporter, and HS10 identifiers. I then aggregate declarations to monthly importer–exporter–variety cells. Throughout the paper, "product" refers to the HS10 code, pooling production origins, and "variety" to the HS10–origin pair; each exercise states which of the two defines its comparison cells. The resulting panel contains roughly 36 million monthly cells and approximately $1.053 trillion in import value. Appendix Table B.3 reports summary statistics for selected years.

Additional Data Sources Exchange-rate data come from Colombia's official historical *Tasa Representativa del Mercado* (TRM) registry. I average the daily TRM to obtain a monthly peso–dollar exchange rate. I use this series to express customs prices in pesos and to separate changes in dollar prices from movements in the exchange rate.

For the pass-through analysis, I combine the TRM with the IMF's monthly exchange rates for suppliers' currencies against the dollar. I assign each foreign supplier to a country and currency using a historical crosswalk that accounts for currency unions and changes in national currencies. I then average the monthly rates to the quarterly level and calculate four-quarter changes.

I use the DANE's official monthly import producer price index to validate the customs-based index. DANE constructs this index from repeated price quotations for specified articles reported by identified importing establishments.[2] As a robustness check, I group the customs observations using the same DANE product categories before constructing the index. This provides a direct comparison between the customs-based and official series.

2.2 Variable Construction

I construct import prices at monthly, quarterly, and annual frequencies by aggregating value and quantity over the relevant calendar months. Let $f \in \{m, q, a\}$ denote the aggregation frequency, and let $\mathcal{M}_f(t)$ be the set of months belonging to calendar period t at frequency f.[3]

For an importer–exporter–variety relationship r, I define the dollar unit value in period t as

$$p_{rt}^{usd,f} = \frac{\sum_{s \in \mathcal{M}_f(t)} V_{rs}}{\sum_{s \in \mathcal{M}_f(t)} Q_{rs}}, \tag{1}$$

where V_{rs} is value and Q_{rs} is quantity in month s. At the monthly frequency, this is simply the value-to-weight ratio. At quarterly and annual frequencies, I first sum value and quantity across the months in the period and then take the ratio.

I construct relationship-level price changes over a k-period horizon as the log difference between current and lagged unit values:

$$g_{rt}^{f,k} = \log p_{rt}^{usd,f} - \log p_{r,t-k}^{usd,f}. \tag{2}$$

A price change is defined only for relationships observed in both $t-k$ and t. I therefore restrict attention to $\mathcal{C}_t^{f,k}$, the set of relationships that continue over the k-period horizon.[4]

I aggregate these relationship-level price changes using lagged import shares. Let $\omega_{r,t-k}^{f,k}$ denote relationship r's share of FOB import value in period $t-k$ among continuing relationships. Dollar-denominated import-price inflation is then given by the geometric Laspeyres index

$$\pi_t^{usd,f,k} = \sum_{r \in \mathcal{C}_t^{f,k}} \omega_{r,t-k}^{f,k} g_{rt}^{f,k}. \tag{3}$$

[2]The quotation record identifies the reporting establishment and article, including its commercial name, brand, internal reference, country of origin, currency, and unit. See DANE's IPP methodology.

[3]Thus, $\mathcal{M}_m(t)$ contains one month, $\mathcal{M}_q(t)$ the three months in quarter t, and $\mathcal{M}_a(t)$ the twelve months in year t.

[4]To limit the influence of changes in unobserved product characteristics within HS10 categories, I follow standard practice and exclude observations with absolute log price changes of three or more.

Let e_t^f denote the average number of pesos per dollar in period t. The corresponding peso index is

$$\pi_t^{cop,f,k} = \Delta_k \log e_t^f + \pi_t^{usd,f,k}.$$

The first term is common to every relationship; the second records changes in dollar-denominated border prices.

Frequency and Horizons I use different frequencies across parts of the analysis. For comparisons with DANE's imported PPI, I construct one- and twelve-month price changes, matching the monthly frequency of the official index. The main accounting exercise instead uses four-quarter price changes measured quarterly. Although both the twelve-month and four-quarter specifications capture inflation over a one-year horizon, the quarterly construction requires a relationship to be observed in the current quarter and in the same quarter one year earlier, rather than in the same month of consecutive years. This reduces sensitivity to shipment timing and increases the share of trade that can be matched across periods. It is also consistent with the approach in Alvarez-Blaser et al. (2025) and Amiti et al. (2024), who measure annual inflation at a subannual frequency. The baseline sample contains 96 overlapping four-quarter observations. After the initial validation in pesos, all price changes and accounting components are expressed in dollars unless a COP superscript is shown explicitly. To simplify notation, I therefore suppress currency, frequency, and horizon superscripts and write the baseline variables as g_{rq}, $\mathcal{C}_q$, $\omega_{r,q-4}$, and π_q.

Sample Coverage While the turnover in importer–exporter relationships is high, continuing relationships account for a disproportionate share of import value, consistent with evidence from other settings (Martin et al., 2023). In the baseline quarter sample, continuing relationships account for 55 percent of current import value on average and no less than 38 percent in any quarter. Coverage is 43 percent for twelve-month changes measured monthly and 57 percent for one-month changes. Thus, relative to exact monthly matching over the same one-year horizon, the quarterly construction materially expands coverage while preserving an annual measure of inflation.

3 Accounting for Imported Inflation

This section first validates the relationship-level customs index and separates peso import-price inflation into its exchange-rate and dollar border-price components. It then studies how concentration across relationships shapes dollar import-price inflation.

3.1 Bottom-Up Imported Inflation

I first validate the relationship-level import-price index against DANE's official imported producer price inflation (IPPI). Panel A of Figure 1 compares the two indexes in pesos. They comove closely and capture the same inflationary and deflationary episodes. Their correlation is nearly 0.9 for twelve-month inflation and 0.8 for one-month inflation.

Panel B of Figure 1 decomposes the customs-based peso index into dollar price inflation and the peso–dollar exchange-rate change. Most variation in peso import-price inflation comes from the exchange rate: in a covariance decomposition, exchange-rate movements account for 92 percent of the variance of the peso index and dollar prices for the remaining 8 percent. This does not make dollar price adjustment economically unimportant, because the two components are negatively correlated. Their correlation is -0.48, and they move in opposite directions in 68 percent of months. Dollar border prices therefore frequently offset exchange-rate movements and materially affect the inflation ultimately realized in pesos.

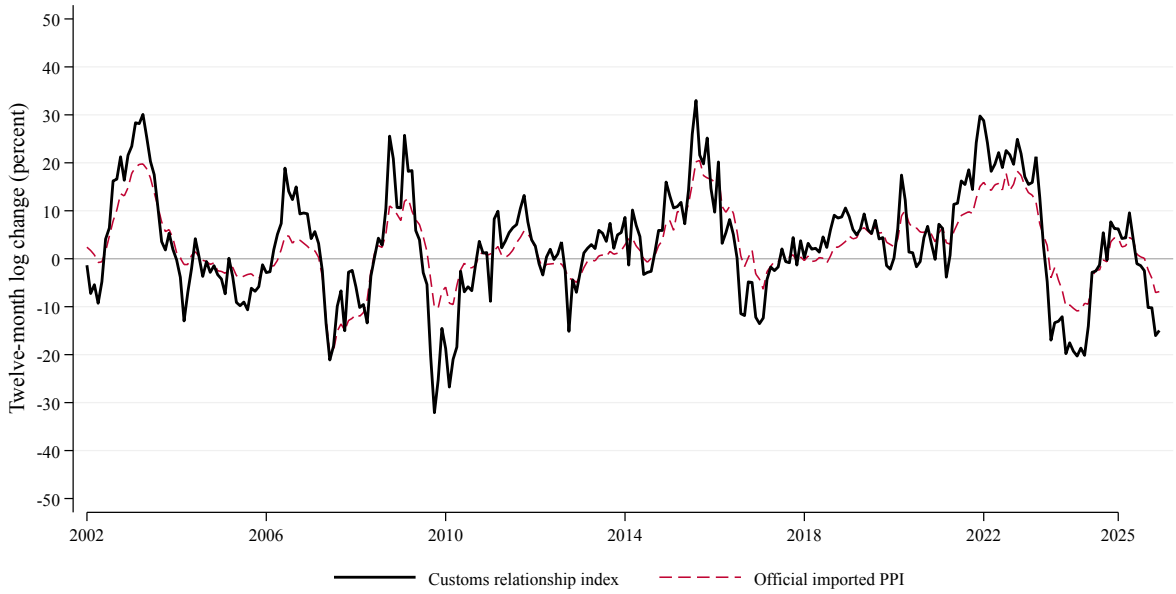
The analysis that follows focuses on dollar price changes because they capture the variation in price adjustment across products and relationships that generates the granular component. The exchange rate enters peso inflation additively and is common to all relationships in a given period. It may affect dollar prices through firms' equilibrium pricing and quantity responses, but it does not mechanically generate differences across relationships and cancels when value- and equal-weighted indexes are compared.

The dollar customs index also contains substantial information beyond the common exchange-rate movement. Its correlation with official peso inflation net of the corresponding exchange-rate change is approximately 0.7 over twelve months (Appendix Table A.1). The relationship-level index therefore captures substantial annual variation in border prices beyond the exchange rate.

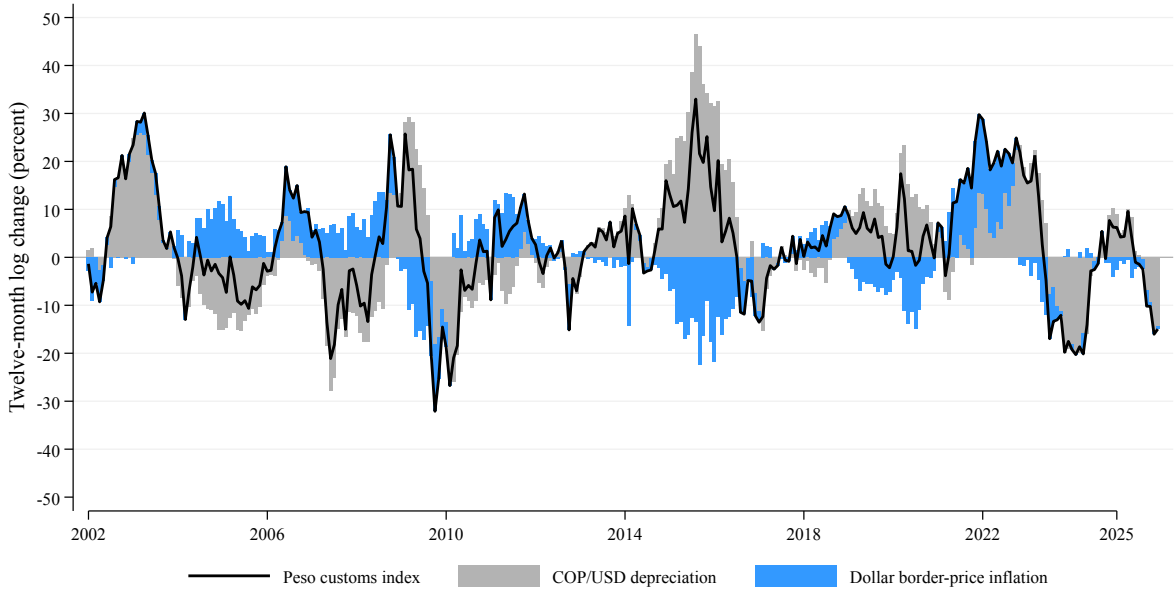
Appendix A reports three additional validation exercises. Neither restricting the relationship index to DANE's product basket nor applying DANE's published weights materially changes its annual correlation with the official series. I also construct conventional product–supplier-country unit-value indexes in the spirit of Amiti et al. (2024), pooling transactions before identifying continuing cells. These indexes track the official series somewhat less closely than the relationship index. When constructed from the same set of continuing relationships, however, the product–supplier-country and relationship indexes are nearly identical. The difference therefore reflects how continuing observations are defined, rather than the level at which prices are aggregated. More broadly, although the relationship index abstracts from entry and exit, continuing relationships capture most aggregate border-price movements.

Figure 1: Peso import-price inflation: validation and currency components

Panel A: Customs index and official imported PPI

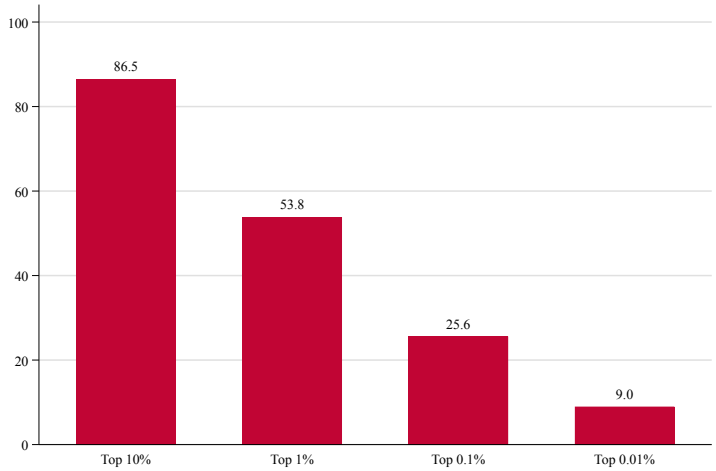


Panel B: Exchange-rate and dollar border-price components



Notes. Twelve-month log changes observed monthly, in percent. In both panels the solid black line is the relationship-level customs index in pesos, constructed from continuing importer–exporter–HS10–standardized production-origin relationships. In Panel A, the dashed line is DANE's official imported producer price index, also in pesos; the correlation between the two series is 0.891. In Panel B, the gray bars report the COP/USD change, with positive values denoting peso depreciation, and the blue bars report dollar border-price inflation. The two components sum exactly to the black line; components with opposite signs extend from zero in opposite directions. Relationship price changes are aggregated using lagged FOB weights; observations with an absolute log change of three or more are excluded. The sample contains 288 monthly observations from January 2002 through December 2025.

Figure 2: Import-Value Shares of the Largest Relationships



Notes. The figure reports time averages of the cumulative lagged import-value shares held by the indicated upper tails of relationships, in percent. Upper-tail groups are defined separately in each quarter using lagged quarterly import value, with the number of relationships in each group rounded up to the nearest integer. The sample contains relationships observed in both $q-4$ and q with a usable four-quarter unit-value change, matching the baseline accounting sample. A relationship is an importer–exporter–HS10–standardized production-origin cell.

3.2 The Role of Large Relationships

A central insight of Gabaix (2011) and the subsequent granularity literature is that shocks to individual economic units need not wash out in the aggregate when activity is highly concentrated. In this setting, import-price inflation can have granular origins if two conditions hold: import expenditure is concentrated in a small number of relationships, and prices in those relationships move differently from prices in smaller ones. I first establish these two conditions in the Colombian data and then quantify the contribution of large relationships to aggregate import-price inflation.

Concentration of Import Expenditure Previous work documents substantial trade concentration across firms (Freund and Pierola, 2015; Bernard et al., 2018) and bilateral links (Eaton et al., 2011). I show that concentration remains pronounced at the finer level of the buyer–supplier–variety relationship. Relationship-level import shares span several orders of magnitude, are highly skewed, and remain stable over time (Appendix Figure B.1). Disaggregation shifts the distribution of shares to the left relative to those of foreign suppliers, Colombian buyers, and HS6 products, but its upper tail overlaps substantially with the tails of these more aggregated units (Appendix Figure B.2).

To quantify concentration, I use the Herfindahl index and its inverse, which measures the

effective number of relationships:

$$H_q = \sum_{r \in \mathcal{C}_q} \omega_{r,q-4}^2, \qquad N_q^{\text{eff}} = H_q^{-1}, \tag{4}$$

where N_q^{eff} is the number of equally sized relationships that would generate the observed Herfindahl index. The baseline quarterly sample contains approximately 81,000 continuing relationships with usable four-quarter price changes, but the effective number averages only 563, with a median of 576. Annual averages range from 302 to 727 and remain small relative to the number of continuing relationships throughout the sample (Appendix Figure B.3).

Figure 2 reports the average cumulative expenditure shares of the upper tail. The largest 10 percent of relationships account for 86 percent of lagged import value, the largest 1 percent account for 54 percent, and the largest 0.01 percent account for 9 percent. The dispersion in relationship size thus translates into substantial and persistent concentration in import expenditure.

Price Changes at Large Relationships The second condition concerns price adjustment: prices in large relationships must move differently from those in smaller relationships. Existing work documents substantial bilateral price dispersion (Fontaine et al., 2020) and shows that price levels and cost shock pass-through vary with buyer and supplier size (Alviarez et al., 2025). Less is known about whether price changes vary systematically with the size of the buyer–supplier–variety relationship itself. To isolate this margin, I compare four-quarter dollar price growth across relationships importing the same HS10 code in the same quarter, pooling production origins, thereby removing price movements common to the product code.

Within each quarter, I rank continuing relationships by lagged import value and regress their price growth on indicators for the upper tail of the size distribution. All specifications include HS10–quarter fixed effects, so the coefficients compare relationships importing the same HS10 code at the same time. Because import expenditure is highly concentrated at the top, I first distinguish the largest 1% of relationships and then separate the largest 0.1% from the remainder of the top percentile.

Table 1 reports the estimates, while Appendix Figure B.4 shows the full gradient across size bins. Over the full sample, prices in the largest 1% of relationships grow 1.5 percentage points less than prices in other relationships importing the same HS10 code in the same quarter. The difference is largest at the very top: relative to relationships below the top percentile, price growth is 1.4 percentage points lower for the remainder of the top percentile

Table 1: Price Changes Within Products, by Relationship Size

	Full sample	2014–16	2021–22	Full sample	2014–16	2021–22
Top 1% relationship	-1.49***	-4.19***	-0.01			
	(0.34)	(1.16)	(0.95)			
Top 1% excluding top 0.1%				-1.40***	-3.85***	0.04
				(0.34)	(1.19)	(0.87)
Top 0.1% relationship				-2.70***	-8.67***	-0.69
				(0.69)	(2.72)	(2.41)
Observations	7,700,749	1,148,216	834,114	7,700,749	1,148,216	834,114

Notes. The table reports regressions of four-quarter dollar unit-value growth, measured in percentage points, using the baseline quarterly matched sample for 2002–2025. Size indicators are defined from the percentile of lagged relationship import value within each quarter. All specifications absorb HS10–quarter fixed effects. HS10–quarter cells containing only one continuing relationship are excluded because they do not identify within-product differences. Observations are unweighted. Standard errors, reported in parentheses, are two-way clustered by HS10 code and quarter. The 2014–16 and 2021–22 specifications contain only 12 and 8 quarter clusters, respectively, so their standard errors should be interpreted as descriptive. Stars denote significance at the 10, 5, and 1 percent levels.

and 2.7 points lower for the largest 0.1%.[5]

The gradient steepens sharply during the 2014–16 depreciation, when the corresponding coefficients are -4.2 percentage points for the largest 1% and -8.7 points for the largest 0.1%. During the 2021–22 inflation surge, both estimates are close to zero and imprecisely estimated. Differential repricing is therefore not a permanent feature of large relationships but depends on the economic environment, a point I return to below. The evidence nevertheless shows that large relationships can reprice differently from smaller ones, and that this difference can be economically meaningful during periods of substantial exchange-rate and border-price movements.

The evidence establishes the second condition for granularity. Import expenditure is concentrated, and the relationships receiving the largest weights often reprice differently from smaller relationships importing the same HS10 code. I next quantify how much this combination of concentration and differential repricing contributes to aggregate import-price inflation.

[5]The leave-one-out estimates in Appendix Figure B.4 show the same pattern. Residual price growth is 0.3 percentage points among relationships below the median, while price growth in the largest 0.1% is 2.3 points lower than in the bottom half.

3.3 The Two Margins of Granular Imported Inflation

The previous section establishes that import expenditure is concentrated and that repricing varies with relationship size. I now quantify the aggregate implications of these two facts. Following the granular accounting approach of Gabaix (2011), I first measure how unequal relationship size contributes to import-price inflation and then distinguish whether this contribution arises from the products sourced by large relationships or from their differential repricing within products.

I define the simple granular residual as

$$\Gamma_q^s = \underbrace{\sum_{r \in \mathcal{C}_q} \omega_{r,q-4} g_{rq}}_{\text{inflation for a dollar of imports}} - \underbrace{\frac{1}{N_q} \sum_{r \in \mathcal{C}_q} g_{rq}}_{\pi_q^{uw}\text{: inflation at the average relationship}} . \tag{5}$$

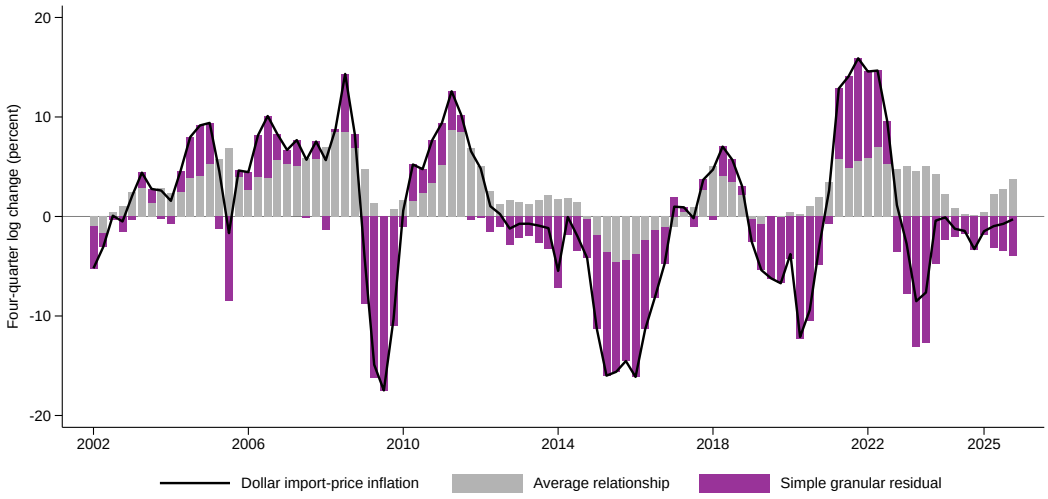
The first term is the baseline value-weighted import-price index, while the second gives equal weight to every continuing relationship. Their difference captures the effect of weighting relationships by their import shares. It is positive when large relationships have higher price growth than the average relationship and negative when they have lower price growth. It is zero if all relationships receive the same weight or experience the same price change.

At the four-quarter horizon, Γ_q^s has a correlation of 0.95 with aggregate dollar import-price inflation, and its covariance with aggregate inflation is equal to 70 percent of the variance of aggregate inflation. Panel A of Figure 3 shows that the residual generally has the same sign as dollar border-price inflation. It is strongly negative during the 2008–09 trade collapse and the 2014–16 decline in dollar import prices, and strongly positive during the 2021–22 global inflation surge. Unequal relationship weights therefore tend to magnify fluctuations in dollar border prices, although the granular residual can be either positive or negative.

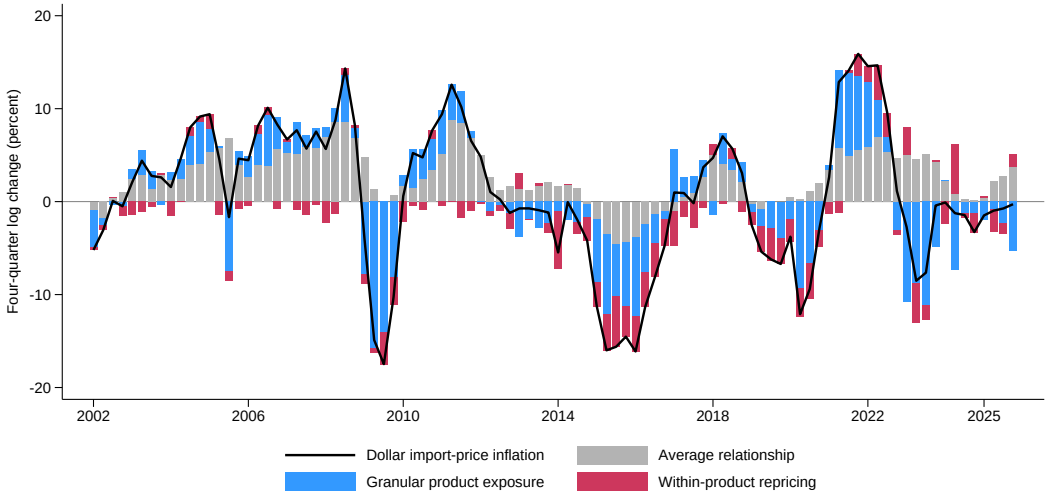
The Origins of the Granular Residual The simple granular residual shows that large relationships shape aggregate inflation, but not why their prices move differently from the average. Two margins can generate the difference: large relationships may be concentrated in products with unusually large common price movements, or they may reprice differently from smaller relationships importing the same HS10 code. I call these *granular product exposure* and *granular repricing*. To separate them, I follow Di Giovanni et al. (2014) and Amiti et al. (2024) and write each relationship's price change as $g_{rq} = \delta_{hq} + u_{rq}$, where δ_{hq} is the unweighted mean price change across relationships importing HS10 code h in quarter q and u_{rq} is relationship r's deviation from it. The common component is defined at the HS10 level, pooling production origins, so that price movements specific to a variety, such as those induced

Figure 3: Dollar import-price inflation and the granular residual

Panel A: The simple granular residual



Panel B: Granular product exposure and within-product repricing



Notes. Four-quarter dollar log changes observed quarterly, in percent. In both panels the black line is the value-weighted dollar import-price index, π_q, and the gray bars are the unweighted price change of the average continuing relationship, π_q^{uw}. In Panel A, the purple bars are the simple granular residual, $\Gamma_q^s = \pi_q - \pi_q^{uw}$; positive bars indicate that large relationships have higher price growth than the average relationship. The residual is unchanged when both indexes are expressed in pesos because the common peso–dollar change cancels in their difference. In Panel B, the blue bars report granular product exposure, G_q^P, and the red bars report within-product repricing, R_q, as defined in equation (6); the two sum to the purple bars in Panel A. In each panel the bars sum exactly to the black line. Components with the same sign are stacked, while components with opposite signs extend from zero in opposite directions.

by the supplier country's exchange rate, remain in the relationship component; Appendix Table B.1 reports the decomposition under an HS10–origin benchmark.[6] Aggregating with lagged import-value weights and subtracting the price change of the average relationship gives

$$\pi_q = \underbrace{\pi_q^{uw}}_{\text{average relationship}} + \underbrace{\sum_r \omega_{r,q-4}\delta_{h(r)q} - \pi_q^{uw}}_{G_q^P\text{: granular product exposure}} + \underbrace{\sum_r \omega_{r,q-4}u_{rq}}_{R_q\text{: within-product repricing}}, \tag{6}$$

with $\Gamma_q^s = G_q^P + R_q$. The exposure term is positive when import expenditure is concentrated in products whose prices rise faster than the average product; the repricing term is positive when large relationships see faster price growth than other relationships importing the same HS10 code. I measure each component's contribution by its covariance with π_q divided by $\text{Var}(\pi_q)$; the three shares sum to one.

Panel B of Figure 3 plots the three components. The black line and the gray bars are the same as in Panel A, and the blue and red bars split the simple granular residual plotted there. Over the 96-quarter sample, the covariance shares are 30 percent for the average relationship, 54 percent for granular product exposure, and 16 percent for within-product repricing. Equivalently, common product movements, $\pi_q^{uw} + G_q^P$, account for 84 percent of the covariance with aggregate inflation and within-product movements for the remaining 16 percent.

Granularity in Two Episodes The two major inflationary episodes in the sample produced high peso import-price inflation for very different reasons. During 2014–16, falling oil and other commodity prices weakened the peso but also reduced the dollar prices of many imports. Average four-quarter peso import-price inflation was 7.3%: a 16.4% depreciation was partly offset by a 9.1% decline in dollar import prices. Concentration strengthened this offset. Inflation in the value-weighted index was 7.6 percentage points lower than in the equal-weighted index, with granular product exposure accounting for 4.3 points of the reduction and within-product repricing for 3.3 points. Large relationships were concentrated in products with the steepest dollar-price declines and, within those products, their prices fell more than those paid by other importers.

During 2021–22, dollar import prices and the exchange rate reinforced each other. Dollar prices rose by 10.7% while the peso depreciated by 6.9%, producing average peso import-price inflation of 17.6%. Concentrationraised inflation by 5.4 percentage points relative to the equal-

[6]Because δ_{hq} includes relationship r's own price change, u_{rq} is mechanically zero in product cells with a single continuing relationship and is shrunk toward zero in small cells. The robustness checks below show that excluding small cells or using leave-one-out product means raises the within-product share only modestly.

weighted benchmark, but the composition differed sharply from the earlier episode. Granular product exposure accounted for 4.4 points, while within-product repricing contributed only 0.9 points. Large relationships were concentrated in the products whose prices rose fastest, but within those products they repriced much like other importers. The concentrated sourcing network thus acted as a buffer in the first episode and an amplifier in the second.

The comparison shows that in this setting concentration is neither inherently inflationary nor deflationary. Its effect depends on how expenditure weights comove with price changes. The positive contribution during 2021–22 parallels the post-COVID retail evidence in Alvarez-Blaser et al. (2025), while the 2014–16 episode shows that the same mechanism can dampen inflation. More broadly, common shocks can coexist with granular transmission. When sourcing weights are unequal, a common vector of product-price movements can generate country-specific inflation. Granularity in this case lies in exposure to the shock rather than in the shock itself.

Granularity Across Product Types The decomposition of the granular residual is especially informative when comparing commodities with differentiated goods. For commodities, one might expect most price variation to come from common product-level shocks, since prices are often tied closely to world markets. If so, relationships importing the same narrowly defined commodity should experience similar price changes, and the relationship-specific component should be small. Differentiated goods provide the opposite benchmark, with more scope for bilateral pricing, product heterogeneity, and contract-specific adjustment.

Table 2 repeats the decomposition within the product groups defined in Section 2, renormalizing weights within each group. Commodities account for 33.2 percent of lagged import value in the matched sample and differentiated goods for 58.0 percent; the remaining 8.8 percent is unclassified.

Table 2: Import-price accounting by product group

Group	Value share	Product exposure	Within product	Simple granular	Corr.
Commodity	33.2	53.1	13.1	66.2	0.55
Differentiated	58.0	26.8	22.0	48.8	0.60
Energy	5.8	57.5	22.2	79.8	0.47
Metals	3.8	37.1	6.3	43.5	0.61
Agriculture	8.5	56.6	1.4	57.9	0.12

Notes. Four-quarter dollar unit-value changes observed quarterly. Products are classified using the conservative Rauch (1999) classification. Value shares are averages in the matched accounting sample; the omitted share is unclassified. Product exposure and within-product relationship variation are the two terms in the within-group simple granular residual and therefore sum to the simple granular covariance share. Energy, metals, and agriculture are subsets of commodities and do not exhaust that category. "Corr." is the correlation between the within-product component and group inflation.

The simple granular residual is larger for commodities. Its covariance share is 66%, compared with 49% for differentiated goods. But the source of granularity differs sharply across the two groups. For commodities, 53 percentage points come from granular product exposure and only 13 from within-product relationship variation. Large commodity relationships therefore matter mainly because they are concentrated in products subject to large common price movements. For differentiated goods, the two margins are much closer: 27 percentage points come from product exposure and 22 from within-product relationship variation.

This pattern fits the usual distinction between standardized and differentiated goods, but it also adds an important qualification. Even for commodities, relationship-level price variation does not disappear. Conditional on the same HS10 code, relationships still contribute meaningfully to aggregate price movements. The relationship component is 22 percent for energy, 6.3 percent for metals, and 1.4 percent for agricultural commodities. Agricultural imports come closest to the textbook case of a common world price. For energy, by contrast, more than one-fifth of its covariance with group inflation remains within HS10 cells.

The energy result is particularly useful for interpreting commodity prices in customs data. A commodity may trade in a global market without generating a single observed border price across relationships. Differences in contract terms, grades, transport arrangements, shipment composition, or bilateral pricing can all produce dispersion within a narrow tariff code. The decomposition does not identify which of these mechanisms matter the most, but it shows that common product shocks do not exhaust measured price variation even in an important commodity sector.

Robustness Exercises Appendix Table B.1 reports the relationship covariance share under three alternative constructions of the benchmark sample and three alternative product

definitions. Restricting the sample to products with at least two or at least five continuing relationships, or using leave-one-out product means, moves the relationship covariance share only modestly. Defining the product market more broadly raises the share to 33.2 percent, because price differences across the HS10 codes within a family are assigned to the residual. The most informative variation is the HS10–origin definition. The baseline common component is defined at the HS10 level, so price movements shared by suppliers of the same HS10 code and origin remain in the relationship component; adding production origin to the common component reduces the relationship share to 7.3 percent. Roughly half of the baseline relationship component therefore reflects movements common to relationships within the same variety, while the remainder reflects variation within those cells. Across all benchmarks, economically meaningful price variation remains at the relationship level after conditioning on narrowly defined product movements.

A final robustness check concerns calendar frequency and the horizon over which price changes are measured. Relationship-level movements may be short-lived and disappear with temporal aggregation. At the same time, higher-frequency estimates may be more sensitive to shipment timing and sparse cells. Appendix Table B.2 reports the covariance decomposition across frequencies and horizons. The relationship share is 32.9 percent for one-month changes measured monthly and 16.0 percent for one-quarter changes measured quarterly. At the annual horizon, the estimates are similar across calendar frequencies: 14.6 percent for twelve-month changes measured monthly, 15.7 percent for four-quarter changes measured quarterly, and 14.9 percent when transactions are aggregated annually. Short-run relationship variation is therefore most visible month to month, while the annual-horizon result does not depend materially on whether the underlying calendar is monthly, quarterly, or annual.

3.4 Sourcing Reconfiguration and Fixed-Share Counterfactuals

The decomposition above measures how concentration in the observed network makes aggregate inflation differ from the price change of the average relationship. It does not measure how changes in the network alter Colombia's exposure to foreign shocks. To study this margin, I compare realized import-price inflation with the inflation that the same price changes would have generated under pre-episode sourcing shares. I use 2013 as the baseline for the 2014–16 depreciation and 2019 as the baseline for the 2021–22 inflation surge. These counterfactuals hold price changes fixed and vary only the expenditure weights.

For each quarter q, let $b(q)$ denote the same calendar quarter in the relevant baseline year.

Matching calendar quarters preserves seasonal sourcing patterns.[7] The 2019 counterfactual spans the 2022 revision of the Harmonized System. I therefore define products using the time-consistent HS families described in Appendix C.2, so that changes in tariff codes are not treated as changes in sourcing. A relationship consists of a Colombian importer, a longitudinal foreign-supplier identifier, an HS family H, and a standardized country of production o. When several HS10 codes map to the same relationship, I aggregate their value and quantity before constructing the quarterly unit value. Results based on raw HS10 codes for 2014–16 provide a robustness check.

Relationship-level price changes are observed only for links that trade at both endpoints of the four-quarter window. I therefore report two complementary counterfactuals. The first fixes shares across product–origin cells, providing broad coverage without imputing prices for inactive relationships. The second fixes relationship shares within a common set of links with observed price changes.

Fixed relationship shares on common support. Let x_{rt} denote the import value of relationship r in quarter t, and define its four-quarter price growth as

$$g_{rq} \equiv (\log p_{rq} - \log p_{r,q-4}) .$$

Let $\mathcal{D}_q$ denote the set of relationships that are active in the baseline quarter $b(q)$ and for which g_{rq} is observed. Within this common support, define the actual and baseline expenditure shares as

$$w^A_{rq} \equiv \frac{x_{r,q-4}}{\sum_{j\in\mathcal{D}_q} x_{j,q-4}}, \qquad w^B_{rq} \equiv \frac{x_{r,b(q)}}{\sum_{j\in\mathcal{D}_q} x_{j,b(q)}}.$$

The relationship-level reconfiguration effect is

$$D^{CS}_q \equiv \sum_{r\in\mathcal{D}_q} w^A_{rq} g_{rq} - \sum_{r\in\mathcal{D}_q} w^B_{rq} g_{rq} = \sum_{r\in\mathcal{D}_q} \left(w^A_{rq} - w^B_{rq}\right) g_{rq}. \tag{7}$$

Both terms in equation (7) use the same relationships and price changes. Thus, D^{CS}_q isolates changes in relative expenditure shares among inherited links that remain price-observable. It does not capture entry or exit. Appendix D decomposes D^{CS}_q into product exposure and within-product repricing using harmonized-family product means. Because the product–origin and relationship counterfactuals use different supports, their estimated gaps should not be added.

[7] Because the index uses four-quarter lagged expenditure shares, the actual and baseline shares coincide by construction during the first four quarters after the 2013 baseline.

Episode-level effects. For an episode E, I report the average quarterly reconfiguration effect,

$$\overline{D}_E \equiv \frac{1}{|\mathcal{Q}_E|} \sum_{q \in \mathcal{Q}_E} D_q,$$

where $\mathcal{Q}_E$ contains the quarters in either 2014Q1–2016Q4 or 2021Q1–2022Q4. These statistics are averages of overlapping four-quarter inflation rates, not cumulative price changes.

The peso–dollar exchange-rate movement is common to all relationships. Let e_q denote its four-quarter log change, so that $g_{rq}^{COP} = g_{rq}^{USD} + e_q$. Because both sets of weights sum to one,

$$D_q^{COP} = \sum_{r \in \mathcal{D}_q} \left(w_{rq}^A - w_{rq}^B\right) \left(g_{rq}^{USD} + e_q\right) = D_q^{USD}.$$

The counterfactual gap is therefore the same in dollars and pesos.

Results. Table 3 reports both counterfactuals. In the product–origin exercise, actual inflation averages −3.54 percentage points during 2014–16, compared with −3.11 under 2013 shares. During 2021–22, the corresponding values are 9.51 and 9.83 percentage points. Sourcing reallocation therefore lowers import-price inflation by 0.43 percentage points in the first episode and by 0.32 points in the second. Both gaps arise almost entirely from reallocation across products rather than across origins within products. Common price support covers 85–89% of both current and baseline expenditure.

Table 3: Sourcing reconfiguration on observed price support

Episode	Counterfactual	Actual	Fixed	Actual − fixed	Actual observed/support	Fixed observed/support
2014–16	Fixed product–origin exposure	−3.540	−3.110	−0.430	87.1	84.7
	Fixed shares among surviving baseline links	−7.661	−6.771	−0.890	45.5	44.4
2021–22	Fixed product–origin exposure	9.512	9.827	−0.315	89.3	88.3
	Fixed shares among surviving baseline links	10.852	11.498	−0.646	50.2	48.7

Notes. Entries are episode averages of quarterly four-quarter dollar log changes, in percentage points. The same actual-minus-fixed gaps apply in pesos. Product–origin rows use common observed harmonized-family–origin price support. Surviving-link rows include only inherited relationships with an observed, trimmed four-quarter price change and therefore exclude entry and exit; their final columns report support relative to actual lagged and baseline value. The two counterfactuals use different supports and their gaps must not be added.

At the relationship level, reallocation points in the same direction in both episodes. Despite their very different price environments, current sourcing weights place more emphasis on relationships with lower price growth than do the inherited weights. This shift lowers import-price inflation by 0.89 percentage points in 2014–16 and by 0.65 points in 2021–22. Most of the reduction occurs within products, suggesting that changes in spending across relationships matter more than changes in the product mix. The exact division between these

margins depends on how the product mean is constructed, but the total gap does not.

These results should nevertheless be interpreted with caution. The common-support sample covers about half of current and baseline expenditure, although it accounts for roughly 76% of the harmonized matched-price sample. The estimates are also sensitive to excluding dominant relationships, and the pre-episode placebo gaps are economically sizable. I therefore view them as descriptive accounting counterfactuals rather than causal estimates of network adjustment. Appendix D reports the component decompositions and quarterly estimates, while Appendix E presents the more assumption-dependent exercise based on imputed prices for the full network.

4 Interpreting Relationship-Specific Repricing

The accounting shows that large relationships reprice differently even within the same HS10 code, but price data alone do not reveal why. A relationship-specific price change may reflect a supplier cost shock, differential pass-through, or a change in the bilateral markup, and price-based decompositions in the inflation literature cannot separate these channels (Alvarez-Blaser et al., 2025; Amiti et al., 2024). The relationship data allow me to go further because the same supplier can be observed selling the same variety to different buyers. This section uses that structure in three steps: a within-supplier decomposition of the granular residual, model-implied bilateral markups, and the pass-through of an observable supplier-country shifter.

4.1 A Stricter Within-Supplier Decomposition

The granular accounting in Section 3.3 defines the repricing component R_q by comparing each relationship with the other relationships importing the same HS10 code, whichever supplier or origin they buy from. I now impose a stricter comparison. I ask whether the *same foreign supplier*, selling the *same variety*, changes its price differently across Colombian buyers, and I use the answer to split the simple granular residual Γ_q^s into a part that is common to all of a supplier's buyers and a part that is specific to individual buyer–supplier relationships.

Following the relationship-pricing framework in Alviarez et al. (2025), I assume that in a given quarter supplier i faces the same change in marginal cost when selling variety h to different Colombian buyers, so that marginal cost varies at the supplier–variety level while markups may also vary across the supplier's bilateral relationships. Four-quarter dollar price

growth for buyer j can then be written as

$$g_{ijhq} = \alpha_{ihq} + \widetilde{\Delta_4 \log \mu_{ijhq}}, \qquad \sum_{j \in \mathcal{J}_{ihq}} \widetilde{\Delta_4 \log \mu_{ijhq}} = 0, \tag{8}$$

where the supplier–variety–quarter effect α_{ihq} collects the change in marginal cost and any markup change common to all of the supplier's Colombian buyers, and the second term is the idiosyncratic markup change of relationship ij relative to the supplier's other buyers. I estimate the idiosyncratic term as the deviation of g_{rq} from the supplier–variety–quarter mean, $\widehat{u}_{rq} = g_{rq} - \bar{g}_{i(r)h(r)q}$. It does not recover the supplier's total markup change: any markup movement shared across buyers is absorbed in α_{ihq} together with costs, so the idiosyncratic term is a lower bound on the role of markups.

Substituting into the definition of the simple granular residual splits it exactly into two parts,

$$\Gamma^s_q = \underbrace{\sum_{r \in \mathcal{C}_q} \left(\omega_{r,q-4} - \tfrac{1}{N_q}\right) \bar{g}_{i(r)h(r)q}}_{\Gamma^{SP}_q\text{: supplier–variety component}} + \underbrace{\sum_{r \in \mathcal{C}_q} \left(\omega_{r,q-4} - \tfrac{1}{N_q}\right) \widehat{u}_{rq}}_{\widetilde{\Gamma^{\mu}_q}\text{: idiosyncratic relationship component}}, \tag{9}$$

the first capturing granular price movements shared across a supplier's buyers and the second those that arise because the supplier prices differently across buyers. As before, either component matters for aggregate inflation only if it differs systematically between large relationships and the average relationship.

Table 4 reports the covariance share of the idiosyncratic component in the granular residual; Appendix Figure B.5 plots the quarterly series. In the full sample the share is 0.5 percent, but this is close to zero partly by construction. When a supplier sells a variety to a single Colombian buyer in a quarter, which is the case for most import value and for many of the largest relationships, there is no cross-buyer variation and the residual is identically zero. The remaining columns therefore restrict attention to supplier–variety–quarter cells with at least two Colombian buyers, which account for 25.1 percent of lagged import value overall, 36.1 percent for commodities, and 19.6 percent for differentiated goods, with weights renormalized within this subsample.

Table 4: Within-supplier decomposition of the simple granular residual

	Full sample	Supplier–varieties with two or more buyers	
	Idiosyncratic share	Value coverage	Idiosyncratic share
All imports	0.5	25.1	3.5
Commodities	0.4	36.1	1.5
Differentiated goods	1.5	19.6	17.2

Notes. Four-quarter dollar unit-value changes observed quarterly, 2002–2025. Idiosyncratic share is the covariance share of the idiosyncratic relationship component, $\Gamma_q^{\widetilde{\mu}}$, in the simple granular residual, $\mathrm{Cov}(\Gamma_q^{\widetilde{\mu}}, \Gamma_q^s)/\mathrm{Var}(\Gamma_q^s)$, in percent; the supplier–variety fixed-effect component, Γ_q^{SP}, accounts for the remainder. Value coverage is the share of lagged import value in supplier–variety–quarter cells observed with at least two Colombian buyers; weights are renormalized within each sample and product group. A variety is an HS10 code paired with its production origin. Commodities and differentiated goods follow the conservative Rauch (1999) classification; unclassified products enter only the all-imports row.

Even within this sample, most granular price movements are shared across a supplier's buyers. The idiosyncratic component accounts for 3.5 percent of the granular residual's covariance overall and 1.5 percent for commodities. For differentiated goods it is substantially larger, at 17.2 percent, with the remaining 82.8 percent captured by the supplier–variety component. The evidence therefore points away from buyer-specific bilateral markup adjustment as the main driver of granular repricing, although it cannot rule out markup movements common to all of a supplier's buyers, which the next exercise addresses.

4.2 Markups Implied by Bilateral Shares

The price comparisons above do not reveal whether repricing shared across a supplier's buyers reflects changes in marginal costs or markups. I separate these components using the two-sided bargaining model of Alviarez et al. (2025). In the model, the price in relationship r is a bilateral price-cost ratio over the supplier's marginal cost,

$$p_{rq} = \mu_{rq} mc_{rq},$$

where the price-cost ratio is a function of two observed market shares:

$$\mu_{rq} = \mu(s_{rq}, x_{rq}; \varphi, \theta, \rho, \eta)\,. \tag{10}$$

Here, s_{rq} is the supplier's share of the buyer's expenditure on the variety, and x_{rq} is the buyer's share of the supplier's sales of that variety. A larger supplier share raises the price-cost ratio, while a larger buyer share lowers it. I calculate both shares each quarter using the

full network of eligible relationships. A variety is an HS10 code paired with its country of production. I take the model parameters from Alviarez et al. (2025).[8] Sensitivity checks not reported here show that the results are robust to alternative parameterizations.

Given the model-implied markup, I recover marginal-cost growth from the identity

$$\Delta_4 \log p_{rq} = \Delta_4 \log \mu_{rq} + \Delta_4 \log mc_{rq}. \tag{11}$$

For relationships in the baseline matched sample, I remove HS10–origin–quarter fixed effects from price, markup, and marginal-cost growth. I then compute the covariance contribution of markups and marginal costs to relationship-specific price changes using lagged import-value weights. By construction, the two covariance shares sum to one.

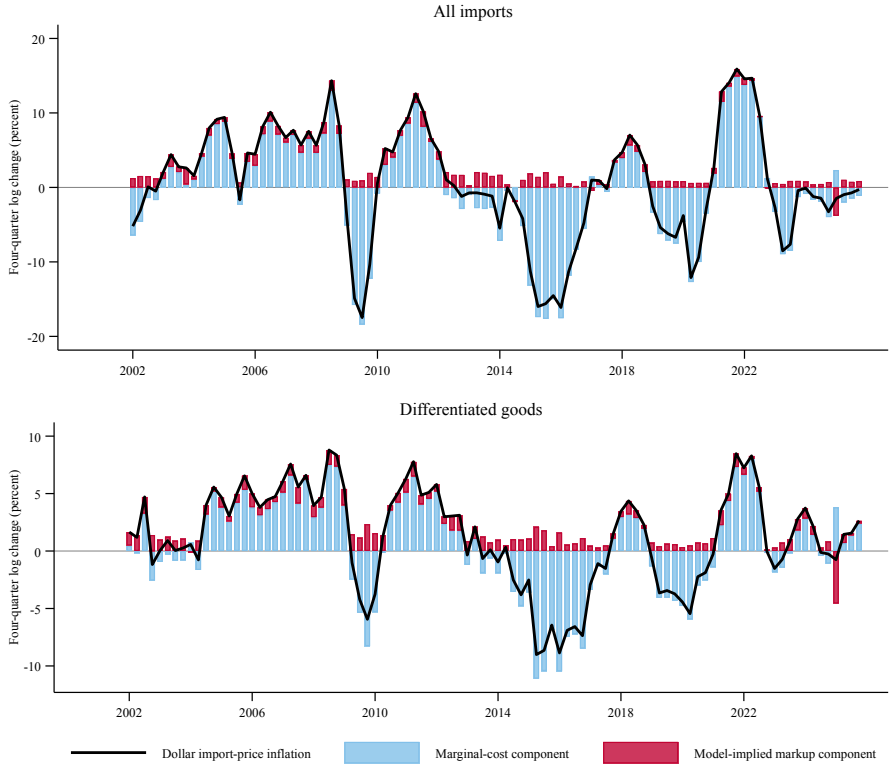
This exercise assigns little of relationship-specific repricing to markups. Markup changes account for only 2.1% of the covariance with price changes across all imports, 2.2% among differentiated goods, and 1.0% among commodities. The estimates are precise: two-way-clustered standard errors are below 0.5 percentage points in every group. The small covariance shares do not arise because model-implied markups are constant. Their variance is 70% as large as the variance of relationship-specific price changes, but the two series are nearly uncorrelated. When I set $\varphi = 0$ and remove buyer bargaining, the markup covariance share falls below 0.3% in every group. Thus, even the small contribution of markups depends on two-sided bargaining rather than supplier concentration alone.

The aggregate decomposition leads to the same conclusion. Figure 4 decomposes four-quarter dollar import-price inflation into model-implied markup and marginal-cost components, using the same lagged import-value weights as the baseline index. Aggregate import-price inflation has a standard deviation of 7.6 percentage points, compared with 0.7 points for the markup component. The markup covariance share is zero to rounding, both for all imports and for differentiated goods. The contrast is especially clear in 2015: dollar import-price inflation averaged -14.4 percentage points, while the markup component contributed $+1.5$ points. The large movements in aggregate import prices are therefore assigned almost entirely to marginal costs rather than to changes in bilateral markups.

Appendix Table B.4 adds a reduced-form check. In depreciation quarters, relationships in the top 0.5 percent of lagged import value cut dollar prices by an additional 3.3 percentage points relative to other relationships in the same HS10–quarter cell. The differential vanishes once Colombian buyers are compared within the same exporter–variety–quarter cell, and

[8] I set the buyer's bargaining weight to $\varphi = 0.845$, the supply parameter to $\theta = 0.45$, and the elasticities in the oligopoly block to $\rho = 10$ within products and $\eta = 2.5$ across products. The model bounds the price-cost ratio between θ and $\eta/(\eta - 1)$, so no trimming of markup changes is required.

Figure 4: Model-implied markup and marginal-cost components of import inflation



Notes. Four-quarter dollar unit-value changes observed quarterly, 2002–2025, on the baseline matched relationship sample. The black line is dollar import-price inflation with lagged import-value weights. Cranberry bars report the model-implied markup component, which aggregates changes in bilateral markups computed from observed quarterly shares with the bargaining model of Alviarez et al. (2025); blue bars report the marginal-cost component, the residual in equation (11), which includes source-country exchange-rate movements. The two components sum exactly to the line. The bottom panel restricts the sample and weights to differentiated goods under the conservative Rauch (1999) classification.

interactions with relationship age, exclusivity, and the bilateral shares that determine model-implied markups are small and leave it unchanged. The repricing of large relationships is therefore a supplier-level phenomenon rather than a change in bilateral bargaining conditions.

Discussion The stability of markups reflects the persistence of market structure in firm-to-firm trade. In the model, bilateral markups depend only on the buyer and supplier shares, s_{rq} and x_{rq}, in equation (10). These shares move slowly. Within varieties, the value-weighted four-quarter correlation of the Herfindahl index is 0.83 for supplier shares and 0.86 for buyer shares, and the value-weighted median absolute four-quarter change is 0.06 for supplier shares and 0.05 for buyer shares. Consistent with this persistence, the value-weighted average log markup implied by these shares moves within a range of 5.7 log points over the sample, and its four-quarter changes have a standard deviation of 0.9 points, compared with 7.6 percentage points for four-quarter dollar import-price inflation. Large movements in border

prices therefore occur with little change in the bilateral shares that determine markups. Exchange rates, world prices, and supplier costs can move transaction prices substantially without materially altering relative purchase and sales shares.

This pattern is consistent with bottom-up models of markup fluctuations (Burstein et al., 2025), in which markups respond to shocks that reallocate market shares. Idiosyncratic shocks can generate such reallocation, whereas shocks that are common within a product may move prices without changing relative shares. In the Colombian data, extensive entry, exit, dormancy, and re-entry coexist with persistent concentration and bilateral shares among continuing links; model-implied markups therefore move little during episodes of large import-price adjustment. The evidence suggests that markup adjustment depends less on the magnitude of the price shock itself than on whether the shock reshapes the distribution of market shares.

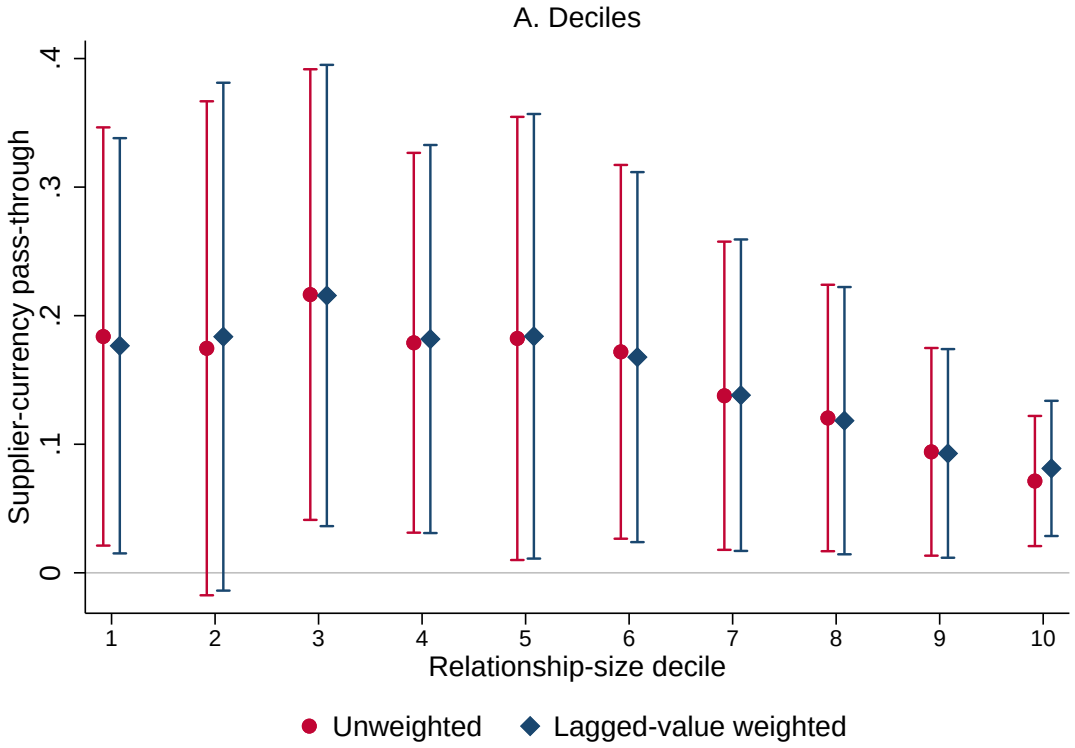
4.3 Heterogeneous Transmission of Supplier-Currency Shocks

The preceding exercises use the cross-sectional structure of the sourcing network to interpret repricing. I now start from an observable foreign shifter and ask whether its transmission into border prices depends on relationship size. If a given supplier-country shock passes through differently across relationships, the network transforms the shock before expenditure weights aggregate it. I use the four-quarter appreciation of the supplier country's currency against the U.S. dollar as a reduced-form shifter of the supplier's dollar costs, attaching the supplier country through the longitudinal exporter identifier, so that it need not coincide with the product's production origin. I rank relationships by lagged import value within each quarter, divide them into ten equally populated groups, and estimate on the exchange-rate-matched sample

$$g_{rq} = \alpha_{hq} + \sum_{d=1}^{10} \lambda_d \mathbf{1}\{D_{rq} = d\} + \sum_{d=1}^{10} \beta_d \mathbf{1}\{D_{rq} = d\} \Delta_4 \log e_{c(r)q} + \varepsilon_{rq}, \tag{12}$$

where $e_{c(r)q}$ is dollars per unit of supplier-country currency, D_{rq} is the relationship-size decile, and α_{hq} is an HS10–quarter fixed effect, so that β_d compares suppliers exposed to different currency movements within the same HS10 code and quarter. The preferred estimates weight observations by lagged relationship value; I also report unweighted estimates. Standard errors are two-way clustered by currency and calendar quarter, the dimensions at which the shifter varies.

Figure 5: Supplier-currency pass-through by relationship size



Notes. The figure reports separate estimates by relationship-size decile of the pass-through of a four-quarter supplier-currency appreciation against the U.S. dollar into four-quarter dollar unit-value growth. Deciles are defined within quarter from lagged relationship value, with decile 1 the smallest relationships. All regressions absorb HS10–quarter fixed effects. Circles are unweighted estimates and diamonds use lagged relationship value weights; vertical bars are 95 percent confidence intervals from standard errors two-way clustered by currency and calendar quarter. The sample covers 2002Q1–2025Q4. Supplier country is assigned through the foreign exporter rather than the country of production.

Figure 5 shows that pass-through is broadly flat across the lower half of the size distribution and declines through its upper half. The value-weighted elasticity falls from 0.177 in the smallest decile to 0.081 in the largest, and the unweighted estimate from 0.184 to 0.071; a pooled specification that places all deciles on common support gives 0.240 and 0.084. Equality of the ten coefficients is strongly rejected ($p < 0.001$). A 10 percent appreciation of the supplier's currency is thus associated with roughly 2.4 percent dollar price growth at relationships in the smallest decile but only 0.8 percent at those in the largest.

These results sharpen the paper's main point. The external disturbance is the vector of supplier-country currency movements. The sourcing network determines which elements of that vector receive large expenditure weights, as in Section 3.4, and also the fraction of each movement that reaches transaction prices. Two economies facing the same foreign shocks, or the same economy under two sourcing configurations, can therefore record different import-price inflation before any domestic response to the aggregate terms of trade.

5 Macro Implications

This section examines whether the composition of import-price inflation matters after aggregation. I study two outcomes: the contemporaneous response of domestic manufacturing to changes in imported-input costs and the subsequent evolution of the border prices underlying those changes. The first captures downstream transmission, while the second captures persistence.

5.1 Downstream sectoral transmission

To study downstream transmission, I map the HS10 product into 35 imported-input groups in Colombia's product-use matrix and then into manufacturing activities using the 2014 input-output structure. This produces, for each activity s and quarter q, a common-product cost shock C^P_{sq} and a relationship-repricing cost shock C^R_{sq}, each scaled by the activity's imported-input share in total intermediate use. Their sum is the total imported-input cost shock, $C^M_{sq} = C^P_{sq} + C^R_{sq}$. I then estimate

$$\Delta_4 y_{sq} = \alpha_s + \lambda_q + \beta_P C^P_{sq} + \beta_R C^R_{sq} + \theta \mu_s \Delta_4 \log E_q + u_{sq}, \tag{13}$$

where y_{sq} is the domestic-output producer-price index, real production, real sales, or employment; E_q is the peso–dollar exchange rate; and μ_s is the imported-input share. Sector fixed effects absorb persistent differences across industries, while quarter fixed effects absorb aggregate Colombian conditions. The remaining variation compares sectors with different product-use exposure to the same quarter's transaction-price movements. Because realized border prices may respond to domestic conditions, the estimates are descriptive correlations rather than causal responses to external shocks.

Table 5 shows that a one-percentage-point imported-input cost increase is associated with approximately 0.56 percentage points of domestic producer-price inflation over four quarters. The estimate is virtually unchanged whether the increase originates in common product prices or in relationship repricing. For production, sales, and employment, neither component nor the difference between them is distinguishable from zero. The low within-sector correlation between the cost components, 0.08–0.09, indicates that the equality is not simply a collinearity artifact.

At this level of aggregation, the evidence is therefore consistent with a producer responding to the cost increase it faces regardless of whether that increase arises from common product-price movements or from relationship-specific repricing. Hence, network structure matters

Table 5: Imported-input cost shocks and domestic manufacturing

Outcome	Common product $\hat{\beta}_P$	Repricing $\hat{\beta}_R$	Difference $\hat{\beta}_P - \hat{\beta}_R$	p: equal	1 s.d. common effect	1 s.d. repricing effect	Total cost $\hat{\beta}_M$	Obs. / sectors
Domestic-output PPI inflation	0.559	0.561	-0.002	0.995	0.964	0.389	0.559	1,320 / 33
	(0.287)	(0.476)	(0.363)				(n.a.)	
	[0.289]	[0.411]	[0.271]	[0.994]			[0.296]	
Real production growth	-0.181	0.179	-0.360	n.a.	-0.302	0.120	-0.121	1,672 / 38
	(n.a.)	(n.a.)	(n.a.)				(0.396)	
	[0.462]	[0.371]	[0.507]	[0.485]			[0.407]	
Real sales growth	-0.193	0.327	-0.520	n.a.	-0.323	0.219	-0.106	1,672 / 38
	(n.a.)	(n.a.)	(n.a.)				(0.376)	
	[0.433]	[0.360]	[0.512]	[0.320]			[0.376]	
Total employment growth	-0.011	0.179	-0.190	0.336	-0.018	0.120	0.021	1,672 / 38
	(0.147)	(0.216)	(0.193)				(0.143)	
	[0.144]	[0.254]	[0.252]	[0.459]			[0.139]	

Notes. The table reports estimates of equation (13). Price components are aggregated from continuing importer–exporter–HS10–origin relationships to MUPNI product groups using lagged import-value weights and mapped into manufacturing activities with the 2014 input-use structure. Outcomes are four-quarter changes; cost shocks are percentage points of total intermediate-input cost. All specifications include activity and quarter fixed effects and the peso depreciation interacted with the activity's imported-input share. Parentheses contain standard errors two-way clustered by MUPNI activity and quarter; brackets contain the activity-clustered estimates used when the two-way covariance matrix is not positive definite. "1 s.d. effect" multiplies each coefficient by its within-fixed-effect standard deviation. The total-cost column imposes $\beta_P = \beta_R$. Coefficients are associations, not causal effects.

for the formation of the imported-input price index, but the current sectoral index appears to summarize its contemporaneous association with domestic producer prices. Aggregation plays an important role in this regard: idiosyncratic relationship repricing is averaged within product groups before it reaches the input-use matrix, so the exercise cannot recover firm-level strategic complementarities or substitution across suppliers. A more compelling test of the network's role in downstream transmission would require importer-level data on domestic production, sales, and employment, which are not available in the current dataset.

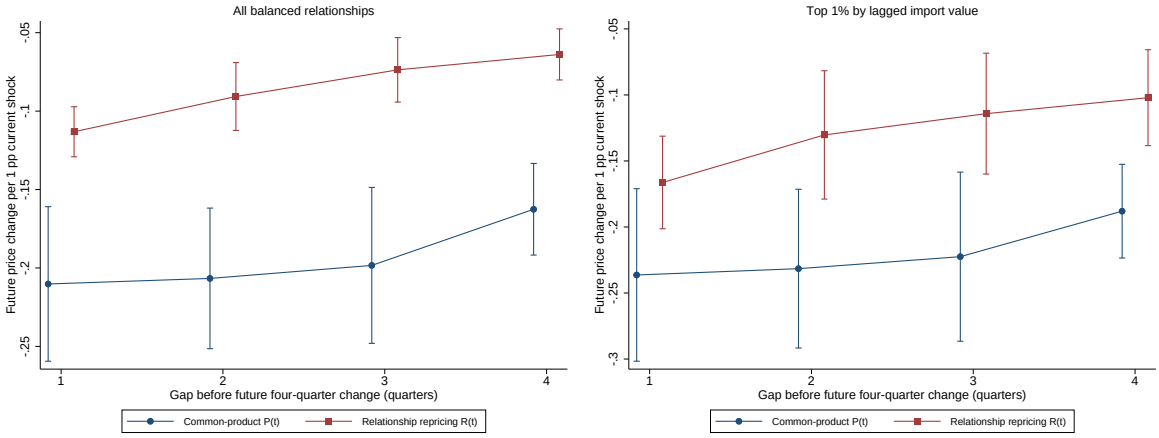
5.2 Different Origins, Different Subsequent Price Paths

Equal contemporaneous transmission does not imply equal dynamics. The usual projection of future four-quarter inflation on current four-quarter inflation cannot answer this question at horizons below one year because the two windows overlap mechanically. It can also manufacture reversal when the same endpoint unit value enters the current change positively and the future change negatively.

I avoid both problems by following the same relationship into periods strictly after t. For relationship r, the current four-quarter price change is

$$x_{rt} \equiv \log p_{rt} - \log p_{r,t-4} = P_{h(r)t} + R_{rt}, \tag{14}$$

Figure 6: The subsequent dynamics of product-price movements and relationship repricing



Notes. The figure reports estimates of the relationship-level projection described in equation (15). The dependent variable is a future four-quarter dollar unit-value change whose window begins after the current four-quarter shock; the two changes share neither quarters nor an endpoint price. The balanced sample requires the exact importer–exporter–HS10–origin relationship at all four future endpoints and holds its $t-4$ import-value weight fixed. All regressions include relationship and quarter fixed effects, use lagged import-value weights, and cluster standard errors by HS10. The right panel restricts the sample to the top one percent of relationships by lagged value in the current quarter. Bars report 95 percent confidence intervals.

where $P_{h(r)t}$ is the common price change for product $h(r)$ and R_{rt} is the relationship-specific component. I compare this initial change with the subsequent four-quarter change

$$y_{rtg} \equiv \log p_{r,t+g+4} - \log p_{r,t+g}, \qquad g = 1, \ldots, 4. \tag{15}$$

The future window begins between one and four quarters after t. It therefore neither overlaps the initial window nor uses p_{rt} as an endpoint.

For each horizon g, I estimate

$$y_{rtg} = \beta_g^P P_{h(r)t} + \beta_g^R R_{rt} + \alpha_r + \delta_t + \varepsilon_{rtg}, \tag{16}$$

where α_r and δ_t are relationship and quarter fixed effects. A negative coefficient indicates that part of the initial price increase is subsequently reversed. The primary sample requires each relationship to have valid price changes at all four future horizons, keeping the sample fixed across g. All regressions use the relationship's import-value weight from $t-4$ and cluster standard errors by HS10 product.

Figure 6 shows that common product-price movements reverse more strongly than relationship-specific repricing. A one-percentage-point increase in the common component

is followed by a 0.16–0.21 percentage-point decline, compared with 0.06–0.11 points for the relationship-specific component. The difference is significant at every horizon and remains present among the largest 1% of relationships.

The balanced sample covers 40% of current import value because it requires relationships to remain observed through all four future windows. Horizon-specific samples raise coverage to 49–59% and yield the same ordering. The estimates therefore describe continuing relationships. Measurement error in the relationship-specific component may also attenuate its coefficient toward zero.

The difference remains visible after aggregation to MUPNI product groups, although the national time series is too short to distinguish the two components precisely. The evidence is therefore strongest at the relationship and product-group levels. Together with the sectoral results, this finding shows that the aggregate import-price index is a useful measure of current cost pressure but an incomplete measure of its persistence. The microeconomic origin of an import-price increase matters less for its contemporaneous domestic transmission than for its subsequent path.

6 Conclusion

This paper studies the microeconomic origins of import-price inflation and their macroeconomic implications. Using transaction-level data for Colombia, I show that import-price inflation is granular. The main channel is product exposure: the largest relationships place disproportionate weight on products experiencing large common price movements. Differential repricing in larger relationships also matters, although it accounts for a smaller share of aggregate fluctuations. Pass-through of supplier-currency movements further declines with relationship size. The sourcing network therefore matters both for which foreign shocks receive the most weight and for how strongly those shocks reach border prices. Depending on the shocks hitting the largest relationships, concentration can either amplify or dampen import-price inflation.

The evidence on the mechanism behind differential repricing points to supplier-level costs rather than changing bilateral markups as the main source of granular repricing. Model-implied markup levels vary substantially across relationships, but markup changes account for only 2.1% of the covariance with relationship-level price changes and essentially none of aggregate import-price inflation. This finding provides one piece of evidence in the debate over the role of rising markups in post-pandemic inflation (Weber and Wasner, 2023; Conlon, 2026). In Colombian import data, changes in bilateral margins played little role, including

during the 2021–22 inflation surge.

The granular origins of imported inflation matter in the aggregate, but the macroeconomic implications appear mainly in persistence. The aggregate import-price index therefore summarizes current imported cost pressure reasonably well, while its microeconomic composition contains additional information about how long that pressure will last. Because domestic outcomes are observed at the sector level, however, aggregation may hide heterogeneous responses across importers. The absence of differential sectoral responses does not rule out important effects within sectors, which future research should explore.

The results also speak to how changes in the sourcing network shape the transmission of foreign shocks. I find that, despite extensive turnover in trading relationships, changes in sourcing altered import-price inflation by less than one percentage point in both 2014–16 and 2021–22. In these episodes, network reconfiguration was not a major source of vulnerability to foreign price shocks. This result need not extend to shocks that shift expenditure across the dominant products and suppliers. When such reallocation occurs, the same foreign price movements can produce a different national import-price index. Counterfactuals in which the sourcing network changes, including through geo-economic fragmentation, must therefore allow import prices to adjust with the network rather than impose a fixed import-price process.

More broadly, the results call for caution when aggregate import prices, and terms-of-trade movements more generally, are treated as primitive shocks in models of small open economies. The observed index combines foreign price movements with country-specific sourcing weights and relationship-level pass-through. It is a useful measure of the cost pressure that reaches the economy, but it is not a complete description of the underlying disturbance. Future work can extend the analysis to export relationships, model the endogenous formation of buyer–supplier networks, and link customs transactions to importer-level prices, quantities, employment, and balance sheets. These extensions would make it possible to distinguish more fully between a foreign shock and the changing trade structure through which the economy experiences it.

References

Alvarez-Blaser, S., R. Auer, S. M. Lein, and A. A. Levchenko (2025): "The Granular Origins of Inflation," NBER Working Paper 33404, National Bureau of Economic Research.

Alviarez, V. I., M. Fioretti, K. Kikkawa, and M. Morlacco (2025): "Two-sided market power in firm-to-firm trade," NBER Working Paper 31253, National Bureau of Economic Research, revised May 2026.

Amiti, M., O. Itskhoki, and J. Konings (2014): "Importers, exporters, and exchange rate disconnect," *The American Economic Review*, 104, 1942–1978.

——— (2019): "International Shocks, Variable Markups, and Domestic Prices," *Review of Economic Studies*, 86, 2356–2402.

Amiti, M., O. Itskhoki, and D. E. Weinstein (2024): "What Drives US Import Price Inflation?" *AEA Papers and Proceedings*, 114, 106–111.

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

Auer, R. A., A. A. Levchenko, and P. Sauré (2019): "International Inflation Spillovers through Input Linkages," *Review of Economics and Statistics*, 101, 507–521.

Berman, N., P. Martin, and T. Mayer (2012): "How do different exporters react to exchange rate changes?" *The Quarterly Journal of Economics*, 127, 437–492.

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

Bricongne, J.-C., L. Fontagné, G. Gaulier, D. Taglioni, and V. Vicard (2012): "Firms and the global crisis: French exports in the turmoil," *Journal of international Economics*, 87, 134–146.

Burstein, A. T., V. M. Carvalho, and B. Grassi (2025): "Bottom-Up Markup Fluctuations," *Quarterly Journal of Economics*, 140, 2619–2684.

Conlon, C. (2026): "Did profits cause inflation?" *International Journal of Industrial Organization*, 105, 103257.

Di Giovanni, J., A. A. Levchenko, and I. Mejean (2014): "Firms, Destinations, and Aggregate Fluctuations," *Econometrica*, 82, 1303–1340.

di Giovanni, J., A. A. Levchenko, and I. Mejean (2022): "Foreign Shocks as Granular Fluctuations," *Working Paper*.

Di Pace, F., L. Juvenal, and I. Petrella (2025): "Terms-of-Trade Shocks Are Not All Alike," *American Economic Journal: Macroeconomics*, 17, 24–64.

Eaton, J., S. Kortum, and F. Kramarz (2011): "An Anatomy of International Trade: Evidence From French Firms," *Econometrica*, 79, 1453–1498.

Fontaine, F., J. Martin, and I. Mejean (2020): "Price Discrimination Within and Across EMU Markets: Evidence from French Exporters," *Journal of International Economics*.

Freund, C. and M. D. Pierola (2015): "Export superstars," *Review of economics and statistics*, 97, 1023–1032.

Gabaix, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79, 733–772.

Gaubert, C. and O. Itskhoki (2021): "Granular comparative advantage," *Journal of Political Economy*, 129, 871–939.

Gopinath, G., E. Boz, C. Casas, F. J. Díez, P.-O. Gourinchas, and M. Plagborg-Mø ller (2020): "Dominant Currency Paradigm," *American Economic Review*, 110, 677–719.

Martin, J., I. Mejean, and M. Parenti (2023): "Relationship stickiness, international trade, and economic uncertainty," *Review of Economics and Statistics*, 1–45.

Mendoza, E. G. (1995): "The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations," *International Economic Review*, 36, 101–137.

Pierce, J. R. and P. K. Schott (2009): "A Concordance between Ten-Digit U.S. Harmonized System Codes and Sic/Naics Product Classes and Industries," *NBER Working Paper, 15548*.

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

Schmitt-Grohé, S. and M. Uribe (2018): "How Important Are Terms-of-Trade Shocks?" *International Economic Review*, 59, 85–111.

Weber, I. M. and E. Wasner (2023): "Sellers' inflation, profits and conflict: Why can large firms hike prices in an emergency?" *Review of Keynesian Economics*, 11, 183–213.

Appendix

A Validation of the Customs Index

This appendix reports the validation exercises summarized in Section 3: the full set of correlations between the customs indexes and DANE's official imported producer price index, a comparison with conventional customs-based indexes, two checks on the continuing-relationship requirement, and a comparison restricted to DANE's product basket.

A.1 Correlations with Official Imported Producer Price Inflation

Table A.1 reports correlations between the customs indexes and official imported PPI inflation at the one- and twelve-month horizons. Columns 1 and 2 compare peso inflation. Columns 3 and 4 ask how much information remains once the common exchange-rate component is removed: they correlate the dollar customs index with official peso inflation net of the corresponding exchange-rate change. For the relationship index, the correlation is 0.098 at the one-month horizon but 0.689 over twelve months. The customs index therefore captures substantial variation in annual border-price inflation beyond the exchange rate, while the weaker monthly correlation may reflect differences in timing between the two indexes.

The comparison remains demanding because the customs and official indexes do not necessarily cover the same products.[9] Section A.4 repeats the exercise for the years in which DANE's product classification and weights are available. Restricting the relationship index to DANE's basket or applying DANE's published weights leaves its annual validation correlations broadly similar. Product matching improves some of the conventional monthly correlations but does not alter the main annual-horizon conclusion.

[9] Alvarez-Blaser et al. (2025), for example, validate their scanner-data price index against an official CPI reconstructed for the product categories represented in their scanner data. In the baseline comparison, I do not restrict the customs index to the products in DANE's collection basket. Subject to the continuing-relationship and trimming requirements, all products in the customs data are eligible to enter the index.

Table A.1: Customs indexes and official imported producer-price inflation

	Official peso inflation		Net of COP/USD change	
Customs index	One month	Twelve months	One month	Twelve months
Continuing firm relationship	0.804	0.891	0.098	0.689
Product–supplier-country indexes on the same relationships				
HS10–supplier country	0.801	0.889	0.154	0.690
HS8–supplier country	0.800	0.890	0.151	0.695
HS6–supplier country	0.789	0.890	0.132	0.694
Conventional product–supplier-country indexes				
HS10–supplier country	0.642	0.851	0.087	0.606
HS8–supplier country	0.636	0.853	0.087	0.608
HS6–supplier country	0.631	0.856	0.058	0.621

Notes. Columns 1–2 report correlations between peso inflation in each customs index and DANE's official imported PPI inflation. Columns 3–4 report correlations between the dollar customs index and official peso inflation net of the corresponding peso–dollar exchange-rate change. Same-relationships indexes use exactly the continuing relationships in the relationship index and then pool their transactions at the indicated product–supplier-country level. Conventional indexes instead pool all transactions before defining unit values and continuing cells. All indexes use lagged FOB weights. The samples contain 299 one-month and 288 twelve-month observations.

A.2 Comparison with Conventional Customs Indexes

Table A.1 also compares the relationship index with conventional customs-based measures. Following Amiti et al. (2024), I construct product–supplier-country unit values at the HS6 level, with HS8 and HS10 as finer alternatives. These indexes differ from the relationship index in both the unit used to match prices over time and the set of transactions that qualify as continuing. To separate these margins, I also construct product–supplier-country indexes using exactly the relationships that enter the relationship index in each period.

The conventional indexes, which pool all transactions before defining continuing product–supplier-country cells, track the official series less closely than the relationship index at every horizon. Their twelve-month peso correlations are 0.851 to 0.856, compared with 0.891 for the relationship index; their one-month correlations are 0.631 to 0.642, compared with 0.804; and their net-of-exchange-rate correlations at twelve months are 0.606 to 0.621, compared with 0.689. The common-support indexes, by contrast, are nearly identical to the relationship index at the twelve-month horizon. Their net-of-exchange-rate correlations with official inflation range from 0.690 to 0.695, and their correlations with the relationship index itself range from 0.995 to 0.996. Once the set of continuing relationships is held fixed, pooling them into product–country cells therefore has little effect on the aggregate index. The weaker performance of the conventional indexes comes from sample composition: cell-level unit values move when the mix of transactions inside a cell changes, and the conventional indexes retain value from relationships that break whenever some other transaction in the cell continues. Matching prices within continuing relationships limits this source of compositional change and preserves the relationship dimension needed for the granular accounting in Section 3.

A.3 Does Continuation Select Unusual Price Observations?

Two further comparisons ask whether the continuation requirement itself selects an unusual set of price observations. Both are run on the baseline quarterly four-quarter sample.

First, I build a conventional HS10–origin index from all eligible transactions, imposing no continuation requirement at the relationship level, so that value from relationships that broke is retained whenever some other transaction in the same cell continues. This index has a correlation of 0.95 with the baseline relationship index and differs from it by 2.1 percentage points in absolute value in the average quarter, with a mean difference of 1.0 point. Part of that gap is the compositional change the relationship index is designed to avoid, since cell-level unit values move when the mix of transactions inside a cell changes. It is therefore an upper bound on what relaxing the continuation requirement alone would do.

Second, I hold the continuing relationships fixed and reweight them so that each HS10–origin cell receives its share of all lagged import value rather than its share of continuing lagged value, with continuers keeping their relative sizes within a cell. This isolates whether continuers sit in an unrepresentative part of the product–origin space. The reweighted index has a correlation of 0.98 with the baseline and differs from it by 1.0 percentage point in absolute value in the average quarter, with a mean difference of 0.01 points, against a baseline standard deviation of 7.6 points. The gap does not systematically favor either index in the major episodes, averaging 0.6 points in 2014–16 and 0.3 points in 2021–22. Cells containing at least one continuing relationship account for 82 percent of the all-import product–origin mix on average, so the exercise cannot speak to cells with no continuer at all. Subject to that limit, conditioning on continuation is not a first-order source of the aggregate movements.

A.4 Matching DANE's Imported-PPI Product Basket

The baseline comparison does not restrict the customs index to the product categories included in DANE's imported PPI. Even with this broader coverage, the twelve-month customs index has a net-of-exchange-rate correlation of about 0.7 with the official series. Following Alvarez-Blaser et al. (2025), I make the comparison more stringent by reconstructing the customs index over a set of product categories that more closely matches the official basket.

This exercise is only possible for the later part of the sample because DANE publishes the relevant imported-PPI classification and fixed category weights beginning in December 2014. I map annual HS10 codes into CPC 2.0 categories using DANE's historical tariff concordance and retain products that can be assigned to the published imported-PPI basket. The mapping covers 290 of 291 categories, accounting for 99.72 percent of DANE's published weight. I then reconstruct the relationship index on this matched sample, using either the baseline customs-value weights or DANE's fixed category weights. For comparison, I also recompute the conventional HS6-, HS8-, and HS10–supplier-country indexes on the same set of products. All results therefore cover January 2015 through December 2025.

Table A.2 mirrors the baseline validation exercise. Its first panel reports the unrestricted indexes over the same post-2014 period, while the remaining panels restrict the customs data to DANE's basket and, in the final specification, apply DANE's category weights. Matching the product basket improves the fit of all customs indexes with the official series. At the twelve-month horizon, the net-of-exchange-rate correlation of the relationship index rises to 0.745, compared with 0.717–0.729 for the conventional product–country

indexes. Using DANE's fixed category weights raises the raw peso correlation of the relationship index to 0.941, the highest in the table, while the net-of-exchange-rate correlation remains 0.732.

Overall, this exercise shows that differences in product composition and weighting between the customs data and DANE's basket do not materially affect the validation of the baseline index. Matching the official basket changes the correlations only modestly.

Table A.2: Customs indexes matched to DANE's imported-PPI product basket

	Official peso inflation		Net of COP/USD change	
Customs index	One month	Twelve months	One month	Twelve months
All customs products, common 2015–2025 period				
Continuing firm relationship	0.874	0.922	0.169	0.749
HS10–supplier country	0.785	0.927	0.122	0.716
HS8–supplier country	0.768	0.927	0.120	0.714
HS6–supplier country	0.756	0.919	0.059	0.718
DANE product basket, customs weights				
Continuing firm relationship	0.882	0.920	0.170	0.745
HS10–supplier country	0.863	0.930	0.194	0.719
HS8–supplier country	0.856	0.930	0.195	0.717
HS6–supplier country	0.844	0.926	0.118	0.729
DANE product basket and CPC weights				
Continuing firm relationship	0.881	0.941	0.168	0.732

Notes. Each entry is a correlation with DANE's official imported-PPI inflation. The first two columns compare peso inflation; the last two compare the dollar customs index with official peso inflation minus the corresponding COP/USD change. All panels cover January 2015–December 2025 and contain 132 observations. The first panel includes all eligible customs products. The second restricts each index to HS10 codes that map into DANE's December-2014 imported-PPI basket and uses lagged customs-value weights. The third first constructs relationship inflation within matched CPC categories and then aggregates those changes using DANE's published fixed weights, renormalized over categories with a valid price change. Conventional indexes pool importers, exporters, and production origins before defining continuing product–supplier-country cells.

B Additional Tables and Figures

Table B.1 examines how the within-product relationship covariance share in equation (6) depends on the construction of the product benchmark and on the definition of the product market. The first rows vary the construction of the product mean. Singleton cells mechanically have zero relationship residuals, and in small cells each relationship contributes substantially to its own benchmark. Restricting the sample to products with at least two or at least five continuing relationships yields relationship covariance shares of 16.7 and 19.1 percent, and using leave-one-out product means gives 18.5 percent, compared with 15.7 percent in the baseline. The remaining rows vary the definition of the product market. The baseline common component is defined at the HS10 level, so price movements shared by suppliers of the same HS10 code and origin remain in the relationship component; defining the common component at the HS10–origin level reduces the

relationship share to 7.3 percent. The time-consistent HS families described in Appendix C.2 link HS6 codes across classification revisions and therefore define broader product groups than HS10. Because these families combine distinct HS10 codes, price differences across them are assigned to the residual: the relationship share rises to 33.2 percent under this definition and to 24.8 percent when origin is also included in the common component.

Table B.1: Relationship-specific covariance shares: definitions and robustness

Definition	Covariance share	Correlation	Periods
HS10 baseline	15.7	0.64	96
HS10, cells with at least two relationships	16.7	0.64	96
HS10, cells with at least five relationships	19.1	0.68	96
HS10, leave-one-out product mean	18.5	0.62	96
HS10 × origin common component	7.3	0.41	96
HS-family robustness	33.2	0.80	96
HS-family × origin common component	24.8	0.69	96

Notes. Four-quarter dollar unit-value changes observed quarterly. Covariance shares are $\mathrm{Cov}(R_q, \pi_q)/\mathrm{Var}(\pi_q)$. The correlation column reports $\mathrm{Corr}(R_q, \pi_q)$. Cell-size restrictions reconstruct both R_q and π_q with weights renormalized within the retained sample. The leave-one-out row removes each relationship from its product mean and uses the corresponding sample of products with multiple relationships. HS families bridge product codes across Harmonized System revisions.

Table B.2 reports the three-part granular accounting at each calendar frequency and change horizon used in the robustness analysis. It shows that the baseline within-product relationship share is similar when yearly inflation is measured from monthly or quarterly observations, while shorter change horizons contain more within-product variation.

Table B.2: Accounting across calendar frequencies and change horizons

Calendar and change horizon	Average rel.	Granular product	Within product	Simple granular	Periods
Monthly, one-month	13.4	53.7	32.9	86.6	299
Quarterly, one-quarter	24.5	59.5	16.0	75.5	99
Monthly, twelve-month	27.3	58.1	14.6	72.7	288
Quarterly, four-quarter	30.2	54.1	15.7	69.8	96
Annual	32.6	52.6	14.9	67.4	24

Notes. Covariance shares of the average-relationship component, granular product exposure, and within-product relationship variation in equation (6). The last two components sum to the simple granular covariance share; all three sum to 100. Monthly one-month and twelve-month changes use lagged monthly values; quarterly one-quarter and four-quarter changes use lagged quarterly values; annual changes use lagged annual values.

Figures B.1 through B.3 complement the concentration evidence in Section 3. The first plots the distribution of relationship import shares by year, the second compares the distribution of import shares across alternative units of aggregation in a representative year, and the third plots the annual average of the quarterly effective number of relationships defined in equation (4).

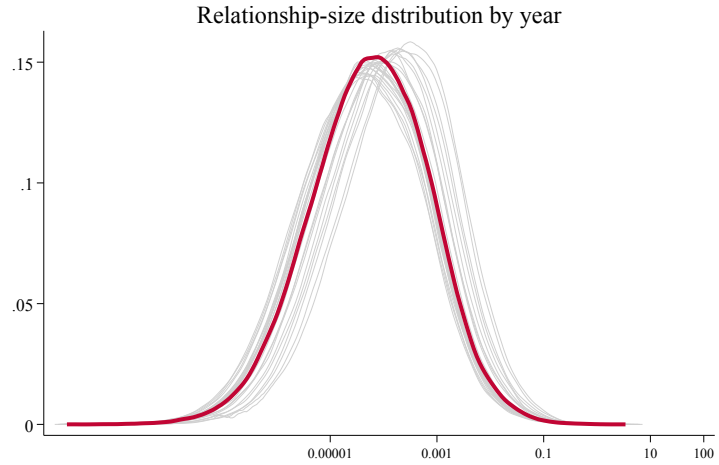
Table B.3 documents the scale of the monthly customs panel at selected points in the sample and reports the median year. It provides the underlying counts behind the concentration measures in the main text.

Table B.3: Scale of the monthly relationship panel

Year	Rel.–months (m)	Relationships	Suppliers	Buyers	HS10	FOB ($bn)
2001	0.70	375,725	66,775	16,015	5,976	11.4
2005	0.95	481,477	65,541	17,986	6,169	19.5
2010	1.26	607,072	75,437	21,590	6,310	37.6
2015	1.62	759,853	90,441	27,079	6,434	51.2
2020	1.46	693,997	84,233	26,516	6,496	40.9
2025	2.04	965,046	103,171	29,511	6,598	66.3
Median year	1.55	734,860	85,875	26,516	6,427	43.3

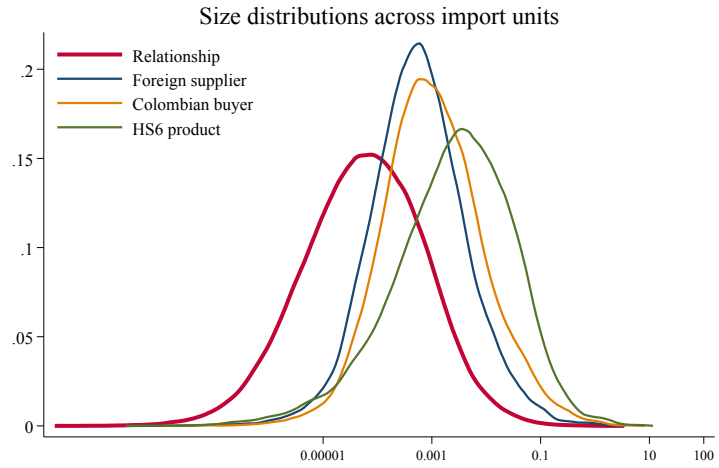
Notes. A relationship is an importer–exporter–HS10–standardized production-origin cell; exporter identifiers are country-specific. Relationship–months count distinct active cells by month. Blank and zero-filled importer identifiers are excluded. FOB values are in current U.S. dollars. The median row is the median across 2001–2025.

Figure B.1: Distribution of relationship import shares, by year



Notes. Kernel densities of lagged quarterly import-value shares across relationships. The horizontal axis reports shares in percent on a log scale; the vertical axis reports kernel density. The figure pools the four quarterly distributions within each year from 2002 through 2025; each gray line represents one year and the cranberry line highlights 2012. The sample in each quarter contains relationships observed in both $q-4$ and q with a usable four-quarter unit-value change, matching the baseline accounting sample. A relationship is an importer–exporter–HS10–standardized production-origin cell.

Figure B.2: Distribution of import shares across alternative import units



Notes. Kernel densities of lagged quarterly import-value shares in 2012. The horizontal axis reports shares in percent on a log scale; the vertical axis reports kernel density. The figure takes the 2012 relationships in Figure B.1 and separately aggregates their lagged values by foreign supplier, Colombian buyer, and HS6 product before recomputing shares within each quarter. A relationship is an importer–exporter–HS10–standardized production-origin cell observed in both $q-4$ and q with a usable four-quarter unit-value change.

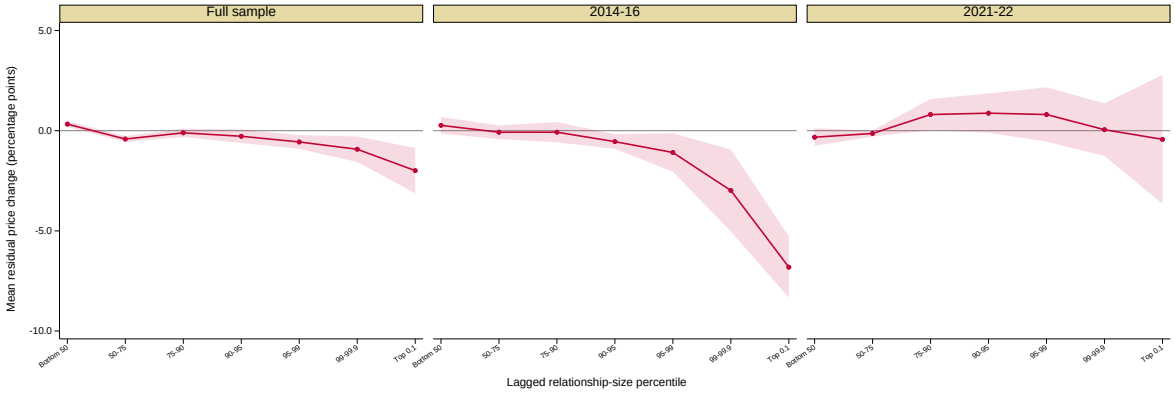
Figure B.3: Effective number of continuing import relationships



Notes. The figure reports the calendar-year average of the quarterly inverse Herfindahl index of lagged relationship-value shares, equation (4); its vertical axis is the effective number of relationships. The underlying quarterly sample contains relationships observed in both $q-4$ and q with a usable four-quarter unit-value change, matching the baseline accounting sample. A relationship is an importer–exporter–HS10–standardized production-origin cell.

Figure B.4 plots the size gradient behind Table 1. For each relationship, I subtract the unweighted mean price change among all *other* continuing relationships importing the same HS10 code in the same quarter. This leave-one-out residual removes common product–quarter movements without allowing a relationship's own price change to affect its benchmark. Relationships are ranked by lagged import value within each quarter, and the top decile is split into progressively narrower groups, ending with the largest 0.1 percent. Over the full sample, relationships below the median have residual price growth of 0.3 percentage points, while the residual declines steadily above the 90th percentile and reaches −2.0 points for the largest 0.1 percent; the difference between the top 0.1 percent and the bottom half is −2.3 points, −7.1 points in 2014–16, and −0.1 points in 2021–22, with a standard error of 1.5 points in the last case. Using the ordinary HS10–quarter residual, which includes each relationship in its own product mean, yields a full-sample difference of −2.0 points.

Figure B.4: Price changes within products, by relationship size: leave-one-out bin means



Notes. The vertical axis reports mean four-quarter dollar price growth relative to other continuing relationships in the same HS10 code and quarter, in percentage points. Size groups are defined from the percentile of lagged relationship import value within each quarter. Points are equal-weighted averages of quarterly bin means. Shaded areas are 95 percent confidence intervals based on Newey–West standard errors with four lags. The full sample contains 96 quarters from 2002q1 through 2025q4; the 2014–16 and 2021–22 panels contain 12 and 8 quarters, so their intervals should be interpreted as descriptive. Relationships in singleton HS10–quarter cells are excluded because a leave-one-out comparison is not defined.

Table B.4 provides a reduced-form test of the mechanism evidence in Section 4. If the differential repricing of large relationships reflects changes in bilateral bargaining conditions, it should covary with observable features of the match, including relationship age, exclusivity, and the bilateral shares that determine model-implied markups. The table regresses four-quarter changes in dollar unit values on the interaction between an indicator for large relationships—those in the top 0.5 percent of lagged import value—and an indicator for quarters in which the peso depreciates by more than ten log points over four quarters. I then add measures of relationship age, exclusivity, and bilateral shares, while progressively tightening the fixed effects from HS10–quarter to HS10–origin–quarter and, for supplier–varieties serving multiple Colombian buyers, exporter–HS10–origin–quarter.

Table B.4: What does relationship size capture?

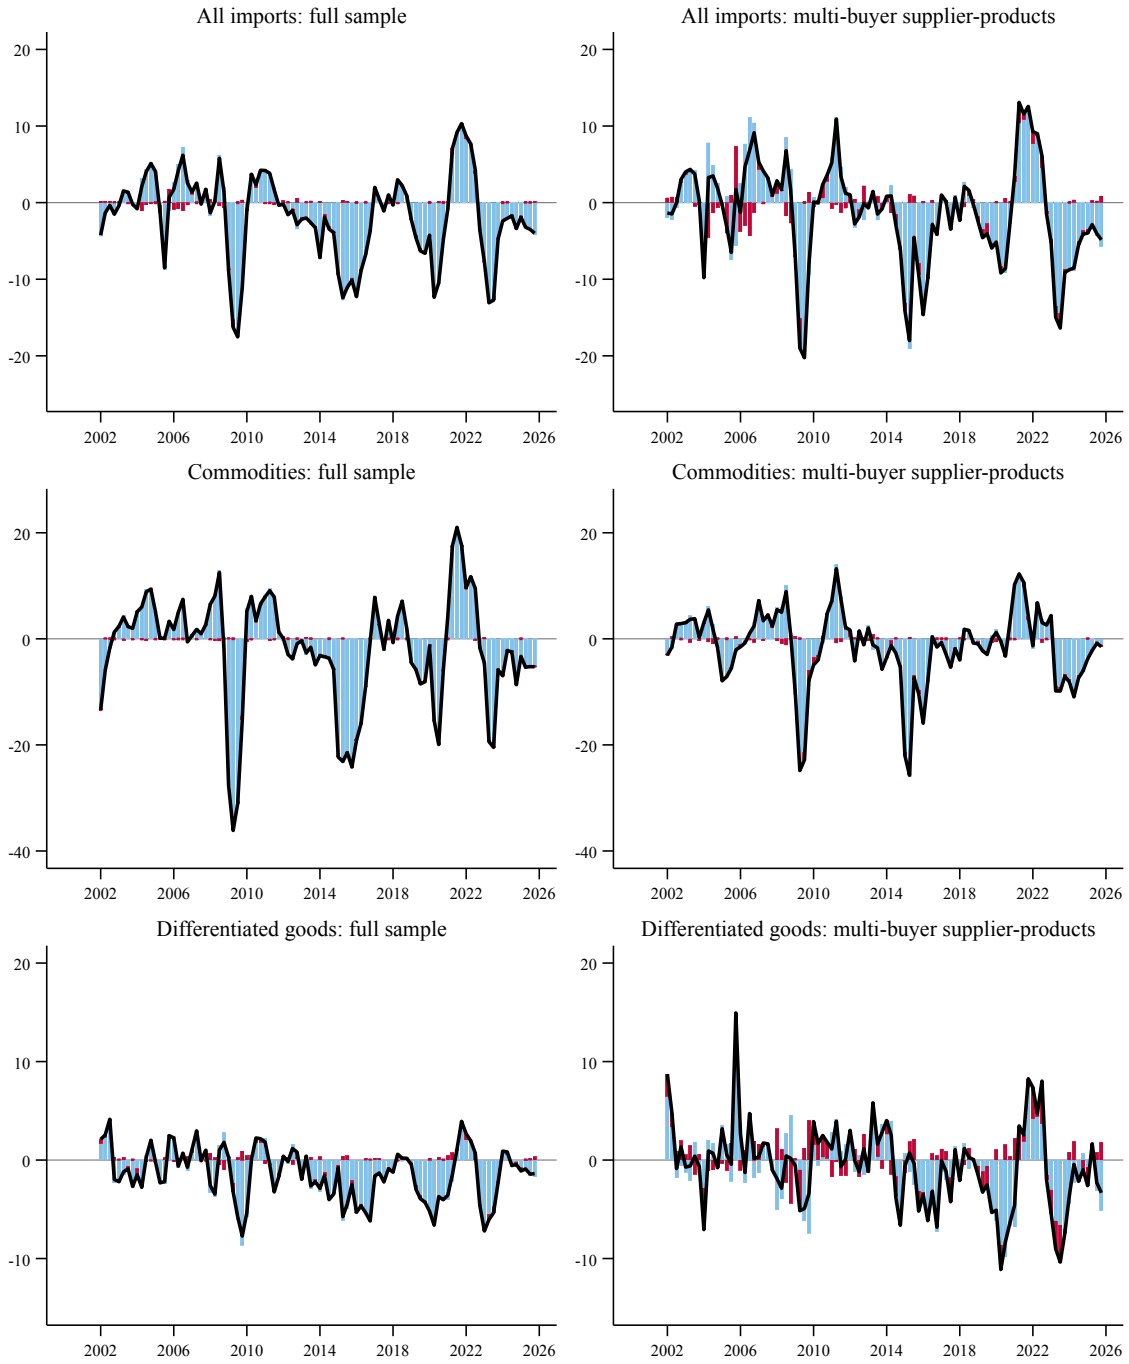
	HS10×q	+origin	Longevity	Exclusivity	Bargaining	Exporter FE
Top 0.5% relationship	-0.0066*	-0.0125**	-0.0136***	-0.0124**	-0.0214***	0.0018
	(0.0039)	(0.0049)	(0.0049)	(0.0049)	(0.0052)	(0.0045)
Top relationship × depreciation	-0.0326***	-0.0216**	-0.0213**	-0.0217**	-0.0207**	-0.0006
	(0.0085)	(0.0098)	(0.0098)	(0.0098)	(0.0098)	(0.0125)
Long relationship × depreciation			-0.0013			
			(0.0034)			
Exclusive supplier–variety × depreciation				-0.0008		
				(0.0043)		
Supplier share × depreciation					0.0026	
					(0.0039)	
Buyer share × depreciation					-0.0023	
					(0.0053)	
Observations	7,700,749	6,672,886	6,672,886	6,672,886	6,672,886	1,242,930

Notes. Regressions of four-quarter changes in relationship-level dollar unit values on the baseline quarterly matched sample, 2002–2025. A depreciation quarter is one in which the peso depreciates by more than ten log points against the dollar over four quarters. "Top 0.5% relationship" indicates the top 0.5 percent of relationships by lagged import value within a quarter. "Long relationship" indicates relationships with trade in at least three prior consecutive calendar years. "Exclusive supplier–variety" indicates that the foreign supplier sells the variety, an HS10 code from a given production origin, to a single Colombian buyer in that quarter. Supplier and buyer shares are the lagged bilateral bargaining shares of equation (10). Columns 1–5 absorb HS10–quarter and then HS10–origin–quarter effects and cluster by HS10 and quarter; column 6 absorbs exporter–HS10–origin–quarter effects on supplier–varieties with at least two Colombian buyers and clusters by exporter–variety and quarter. Stars denote significance at the 10, 5, and 1 percent levels.

During depreciation quarters, large relationships cut dollar prices by an additional 3.3 percentage points relative to other relationships within the same HS10–quarter cell, and by 2.2 percentage points within HS10–origin–quarter cells. This differential vanishes once exporter–HS10–origin–quarter fixed effects are included, which compare Colombian buyers purchasing the same variety from the same foreign supplier in the same quarter. By contrast, interactions with relationship age, exclusivity, and bilateral bargaining shares are small and imprecisely estimated, and their inclusion leaves the size differential essentially unchanged. The evidence therefore attributes the repricing of large relationships primarily to supplier-level forces rather than to changes in bilateral bargaining conditions.

Figure B.5 plots the quarterly series behind Table 4: the simple granular residual and its supplier–variety and idiosyncratic components, for all imports, commodities, and differentiated goods, in the full sample and in supplier–variety cells with at least two Colombian buyers.

Figure B.5: Within-supplier decomposition of the simple granular residual: quarterly series



Notes. Four-quarter dollar unit-value changes observed quarterly, 2002–2025. The black line is the simple granular residual. Blue bars report the supplier–variety fixed-effect component, Γ_q^{SP}, which absorbs marginal-cost changes and markup changes common across a supplier's buyers. Cranberry bars report the residual, $\Gamma_q^{\widetilde{\mu}}$, interpreted under the maintained restriction as idiosyncratic bilateral markup changes; the two bars sum exactly to the line. A variety is an HS10 code paired with its production origin. Rows report all imports, commodities, and differentiated goods. The left column uses the full sample. The right column retains supplier–variety–quarter cells observed with at least two Colombian buyers. Weights are renormalized within each panel. Commodities and differentiated goods follow the conservative Rauch (1999) classification; unclassified products enter only the all-imports panels.

C Data Appendix

This appendix records the data-construction details behind the relationship-level panel: the longitudinal exporter identifier that defines continuing relationships, the country-code and product-code harmonization, and the exchange-rate series used to convert dollar customs prices into pesos.

C.1 Constructing a Longitudinal Exporter Identifier

Colombian importers are identified by their tax number (NIT), but foreign exporters appear as free-text names truncated at forty characters. I therefore construct a country-specific name-continuity identifier after pooling all 300 monthly files. The identifier follows the same exporter through spelling, punctuation, legal-form, and truncation variants. It is not a global corporate-group identifier and does not join mergers, acquisitions, or otherwise unrelated legal names.

Country blocks are constructed from the original customs codes rather than from the displayed Spanish country names. The data use ALADI codes through 2024 and UN M49 codes beginning in 2025; I connect the systems with DANE's official concordance, map ordinary countries to ISO codes, and retain DIAN free zones as separate customs areas. Names are upper-cased and stripped of punctuation and legal-form terms. Exact normalized names form the initial units. Within a standardized supplier-country block, variants are joined using rules designed for the source field: long prefix matches capture forty-character truncation, while conservative Jaro–Winkler and token-sort thresholds capture minor spelling and word-order differences. Short and one-token names face stricter thresholds. Placeholder exporter names receive no identifier and are excluded from the relationship sample.

The resulting crosswalk contains 967,310 normalized country–name units and 785,906 longitudinal exporter identifiers. Of these, 611,944 preserve a strict one-to-one identifier from the preceding crosswalk; split, joined, and newly recovered entities receive new versioned identifiers. No identifier spans standardized supplier countries. The rebuilt monthly panels contain 36,417,719 importer–exporter–HS10–origin cells and $1.053 trillion in FOB value. Adjacent-year FOB linkage averages 89.9 percent at the exporter level and 81.9 percent at the importer–exporter level. Because name linkage has no external ground truth, I use reproducible manual separation rules for identified false joins and retain an audit queue of the 300 highest-value clusters for review.

C.2 Product Classification across HS Revisions

The 2001–2025 sample spans six editions of the Harmonized System (HS 1996, 2002, 2007, 2012, 2017, and 2022) and five revision transitions. A price change computed across a revision boundary on a recoded product would be lost or contaminated. Following the synthetic-product approach of Pierce and Schott (2009), I link any HS6 codes that ever map into one another in the official UNSD correspondence tables and use the resulting connected components—5,000 families covering 6,569 HS6 codes—as time-consistent product identifiers for across-period comparisons; within-period market definitions use raw HS10 codes. The revision timing is verified in the data: codes unique to each edition first carry import value in 2002, 2007, 2012, 2017, and 2022.

C.3 Exchange-Rate Series

Exchange rates are monthly averages of the official peso–dollar market rate (*tasa representativa del mercado*, TRM). All conversions of customs prices into pesos use this common rate. The supplier-currency exercise in Section 4.3 separately uses IMF period-average exchange rates against the dollar for the foreign exporter's country. The quarterly supplier-currency series requires all three monthly observations, measures four-quarter changes against the same currency used in the current quarter, and leaves missing changes unimputed. The sample spans several large, peso-wide swings. The peso depreciated through the early 2000s to about 2,950 per dollar in early 2003, then appreciated through the commodity boom to a trough near 1,770 in early 2012. The collapse in oil prices from late 2014—oil is roughly half of Colombian exports—reversed this sharply: the peso lost about 37 percent against the dollar between 2014 and 2015 and roughly 90 percent from the 2012 trough through early 2016. After stabilizing near 2,950 in 2017–2018, it weakened through the COVID-19 shock (near 3,990 in April 2020) and again in the global tightening of 2022, when the monthly average peaked near 4,930 in late 2022 before easing below 3,800 by the end of 2025.

D Fixed-Share Counterfactuals: Derivations and Additional Results

This appendix records the construction behind the two observed-price counterfactuals in Section 3.4, the decomposition of each gap, and the quarterly series underlying Table 3. Notation follows the main text: q is a quarter in an episode, $b(q)$ the corresponding calendar quarter of the baseline year, H a time-consistent HS family, and o a standardized country of production.

D.1 Fixed Product–Origin Exposure

Let $c = (H, o)$ denote a harmonized product–origin cell and let

$$\bar{g}_{cq} = \frac{1}{N_{cq}} \sum_{r \in \mathcal{C}_q(c)} g_{rq} \tag{D.1}$$

be the unweighted mean four-quarter dollar price change among continuing relationships in that cell. The unweighted mean is consistent with the definition of common product movements in Section 3 and does not embed current relationship size in the cell-level shock. Results using lagged-value-weighted cell means are reported below as a robustness check.

Let $X^{\text{all}}_{c,q-4}$ be the lagged FOB value of all eligible imports in cell c, including transactions outside the continuing-relationship sample, and let $X^{\text{all}}_{c,b(q)}$ be the corresponding value in the baseline quarter. Define $\mathcal{K}_q$ as the set of cells with an observed $\bar{g}_{cq}$ and positive expenditure in both baskets. On this common price support, the actual and fixed shares are

$$S^A_{c,q-4} = \frac{X^{\text{all}}_{c,q-4}}{\sum_{d \in \mathcal{K}_q} X^{\text{all}}_{d,q-4}}, \qquad S^F_{c,b(q)} = \frac{X^{\text{all}}_{c,b(q)}}{\sum_{d \in \mathcal{K}_q} X^{\text{all}}_{d,b(q)}}. \tag{D.2}$$

The actual and fixed indexes and their difference, the product–origin reconfiguration effect used in Section 3.4, are

$$\pi^{PO,A}_q = \sum_{c \in \mathcal{K}_q} S^A_{c,q-4} \bar{g}_{cq}, \quad \pi^{PO,F}_q = \sum_{c \in \mathcal{K}_q} S^F_{c,b(q)} \bar{g}_{cq}, \quad D^{PO}_q \equiv \pi^{PO,A}_q - \pi^{PO,F}_q = \sum_{c \in \mathcal{K}_q} \left(S^A_{c,q-4} - S^F_{c,b(q)} \right) \bar{g}_{cq}. \tag{D.3}$$

I report separately the shares of current and baseline import expenditure covered by $\mathcal{K}_q$.

The product–origin effect can be separated into reallocation across products and across origins within products. Let δ_{Hq} be the unweighted price movement common to harmonized product family H, and write

$$\bar{g}_{Hoq} = \delta_{Hq} + (\bar{g}_{Hoq} - \delta_{Hq}) . \tag{D.4}$$

Substitution into equation (D.3) yields

$$
\begin{aligned}
D_q^{PO} &= D_q^P + D_q^{O|P}, \\
D_q^P &= \sum_{H,o} \Delta S_{Ho,q} \delta_{Hq}, \\
D_q^{O|P} &= \sum_{H,o} \Delta S_{Ho,q} \left(\bar{g}_{Hoq} - \delta_{Hq} \right),
\end{aligned}
\tag{D.5}
$$

where $\Delta S_{Ho,q} = S^A_{Ho,q-4} - S^F_{Ho,b(q)}$. The first term captures changes in the product mix; the second captures changes in origins within products.

D.2 Fixed Relationship Shares on Common Support

Let $\mathcal{B}_{b(q)}$ be the relationships active in the baseline quarter and recall that $\mathcal{C}_q$ contains relationships with a valid four-quarter price change in q. Their common support is

$$
\mathcal{D}_q = \mathcal{B}_{b(q)} \cap \mathcal{C}_q. \tag{D.6}
$$

For relationships in $\mathcal{D}_q$, the actual lagged and fixed baseline weights in equation (7) are

$$
a_{rq} = \frac{V_{r,q-4}}{\sum_{s \in \mathcal{D}_q} V_{s,q-4}}, \qquad b_{rq} = \frac{V_{r,b(q)}}{\sum_{s \in \mathcal{D}_q} V_{s,b(q)}}. \tag{D.7}
$$

I report the shares of baseline expenditure, current lagged expenditure, and matched-sample expenditure represented by $\mathcal{D}_q$.

Using the same observations and weight differences, I further write

$$
g_{rq} = \delta_{H(r)q} + u^H_{rq}, \qquad u^H_{rq} = g_{rq} - \delta_{H(r)q}. \tag{D.8}
$$

Equation (7) then decomposes exactly into product exposure, $\sum_r (a_{rq} - b_{rq}) \delta_{H(r)q}$, and within-product relationship repricing, $\sum_r (a_{rq} - b_{rq}) u^H_{rq}$. Both pieces are constructed entirely on $\mathcal{D}_q$. This harmonized-family split is not the paper's raw-HS10 54.1/15.7 covariance decomposition. The broad product–origin counterfactual in equation (D.3) uses a different support and therefore cannot be added to the relationship estimate.

D.3 Episode Averages

For method $m \in \{PO, CS\}$, episode e is summarized by the average difference in four-quarter inflation,

$$
\overline{D}^m_e = \frac{1}{|\mathcal{Q}_e|} \sum_{q \in \mathcal{Q}_e} D^m_q, \tag{D.9}
$$

where $\mathcal{Q}_e$ contains 2014Q1–2016Q4 or 2021Q1–2022Q4. These are averages of overlapping four-quarter

inflation rates, not cumulative changes.

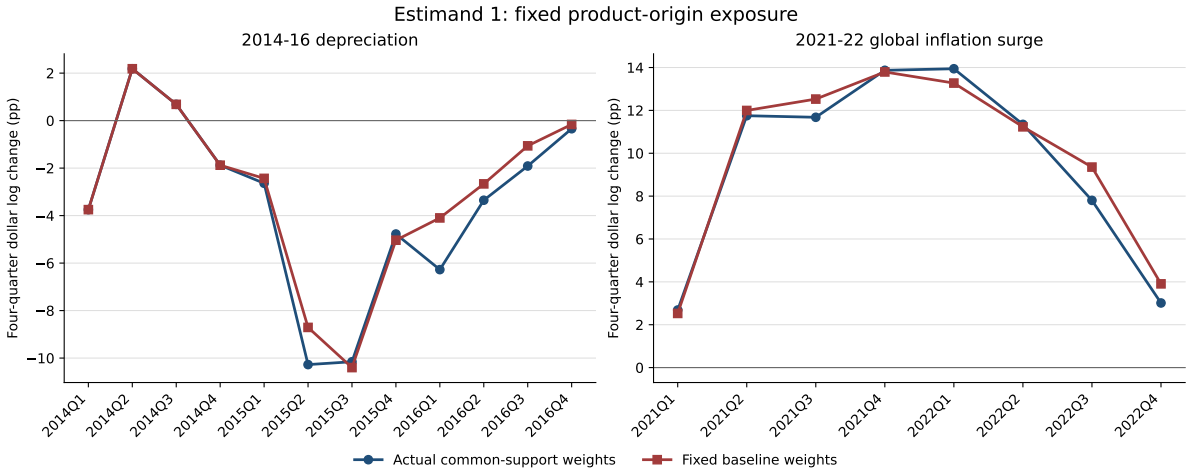
D.4 Component Decompositions and Quarterly Series

For the product–origin counterfactual, the −0.430-point gap in 2014–16 combines a product component D^P of −0.457 points with an origin-within-product component $D^{O|P}$ of 0.027 points. Common price support covers 87.1 percent of current lagged expenditure and 84.7 percent of baseline expenditure. For 2021–22, the −0.315-point gap comprises a −0.377 product component and a 0.062 origin-within-product component, with current and baseline coverage of 89.3 and 88.3 percent.

For the common-support counterfactual, the −0.890-point gap in 2014–16 combines a 0.091 product-exposure contribution with −0.982 points of within-product relationship repricing; the −0.646-point gap in 2021–22 combines −0.236 and −0.410 points. This harmonized-family split is sensitive to how the product mean is weighted: under lagged-value product weighting, the 2014–16 product component becomes −0.540 points and the 2021–22 relationship component becomes 0.370 points, although each total gap is unchanged. The common support represents 45.5 and 50.2 percent of all current lagged value, 44.4 and 48.7 percent of the baseline network, and approximately 76 percent of the harmonized matched-price sample.

Figure D.1 plots the quarterly actual and fixed product–origin indexes in each episode. The two series track each other closely, which is the visual counterpart of the small episode-average gaps in Table 3.

Figure D.1: Reconfiguration relative to inherited shares



Notes. The panels plot quarterly four-quarter dollar import-price inflation under actual common-support product–origin weights and fixed corresponding-quarter baseline weights. The left panel compares actual weights with 2013 shares during 2014–16; the right compares actual weights with 2019 shares during 2021–22. Both series apply the same observed harmonized-family–origin price movements. Their difference is reconfiguration relative to inherited shares; a positive gap means that reconfiguration raised import-price inflation.

E Imputed Full Fixed-Network Counterfactual

The preferred counterfactuals in Section 3.4 avoid assigning bilateral prices to inactive links. A relationship present in a basket but absent at the end of the price-change window has no observed transaction price, and carrying forward its last unit value or assigning it zero price growth would impose an arbitrary counterfactual. This appendix assesses the missing extensive margin by constructing a full-network extension under explicit, testable imputation rules. Because the result depends on many unobserved prices and the prediction diagnostics are weak, it is a sensitivity exercise rather than a headline estimate.

Let $\mathcal{A}_{q-4}$ contain every eligible relationship with positive value in the actual lagged basket, whether or not it survives to q, and let $\mathcal{B}_{b(q)}$ contain every relationship in the baseline basket. Define their union as

$$\mathcal{U}_q = \mathcal{A}_{q-4} \cup \mathcal{B}_{b(q)}. \tag{E.1}$$

For each relationship in $\mathcal{U}_q$, I impute the current four-quarter price growth rate, not a price level. For a link observed in $q-4$ but not in q, the assignment implies $p^*_{rq} = p_{r,q-4}\exp(g^*_{rq})$. For a baseline link absent from both endpoints of the current window, the data do not identify a current price level; the counterfactual instead applies an imputed contemporaneous growth rate to that link's fixed baseline expenditure weight. The resulting object is a fixed-weight shift-share counterfactual, not a reconstructed path of transaction prices.

I construct g^*_{rq} using the following eight-rung hierarchy: (i) the observed harmonized relationship price change, when available; (ii) the leave-one-out mean for the same supplier, raw HS10, and origin across other Colombian buyers; (iii) the raw-HS10–origin mean; (iv) the harmonized-family–origin mean; (v) the raw-HS10 mean; (vi) the harmonized-family mean; (vii) the mean for one of 21 standard HS sections; and (viii) the aggregate equal-weight movement. HS10-specific rungs use the relationship's largest-value raw code, with equal-value ties resolved by the lexicographically smallest code. When a bilateral price is missing, this procedure sets its relationship-specific residual to zero at the narrowest level observable in that quarter. The first fallback is motivated by the within-supplier evidence that most differential repricing is common to a supplier–variety. It is nevertheless an assumption, particularly because supplier–varietys with multiple Colombian buyers account for only 25.1 percent of lagged value and because a relationship's disappearance may be related to its latent price.

Crucially, I use the same assignment g^*_{rq} to price both baskets. Define

$$a^{\text{all}}_{rq} = \frac{V_{r,q-4}}{\sum_{s\in\mathcal{A}_{q-4}} V_{s,q-4}}, \qquad b^{\text{all}}_{rq} = \frac{V_{r,b(q)}}{\sum_{s\in\mathcal{B}_{b(q)}} V_{s,b(q)}}, \tag{E.2}$$

with a weight set to zero when the relationship is absent from the relevant basket. The imputed actual and fixed indexes are

$$\pi^{I,A}_q = \sum_{r\in\mathcal{U}_q} a^{\text{all}}_{rq} g^*_{rq}, \qquad \pi^{I,F}_q = \sum_{r\in\mathcal{U}_q} b^{\text{all}}_{rq} g^*_{rq},$$
$$D^I_q \equiv \pi^{I,A}_q - \pi^{I,F}_q = \sum_{r\in\mathcal{U}_q} \left(a^{\text{all}}_{rq} - b^{\text{all}}_{rq}\right) g^*_{rq}. \tag{E.3}$$

Using the same imputed price function on both sides avoids comparing a fully priced baseline basket with the continuer-only actual index. It also gives an exact decomposition:

$$D_q^I = \sum_{r \in \mathcal{A}_{q-4} \cap \mathcal{B}_{b(q)}} \left(a_{rq}^{\text{all}} - b_{rq}^{\text{all}}\right) g_{rq}^* + \sum_{r \in \mathcal{A}_{q-4} \setminus \mathcal{B}_{b(q)}} a_{rq}^{\text{all}} g_{rq}^* - \sum_{r \in \mathcal{B}_{b(q)} \setminus \mathcal{A}_{q-4}} b_{rq}^{\text{all}} g_{rq}^*. \tag{E.4}$$

The three terms capture reallocation among links in both baskets, links absent from the baseline basket, and baseline links no longer active in the current lagged basket. These labels describe basket membership rather than necessarily first entry or permanent exit; I use each link's complete history to distinguish entry, exit, dormancy, and re-entry.

I assess the price assignment by masking observed bilateral changes and predicting them with the same leave-one-out hierarchy. I report prediction errors and the expenditure share priced at each rung. I also replace missing bilateral changes with product–origin and product-level movements directly and vary the missing relationship residual over observed within-cell quantiles.

Results and diagnostics. Table E.1 reports the imputed rows removed from the main-text counterfactual table. Pricing the union of the actual and baseline networks with the eight-rung hierarchy gives an imputed full-network gap of −1.415 percentage points in 2014–16: the actual all-network index is −7.839 and the fixed index is −6.424. For 2021–22, the corresponding indexes are 10.338 and 11.421, a gap of −1.083 points. These estimates require imputed price growth for 40.1 and 55.6 percent of actual and fixed value in the first episode and 33.6 and 51.3 percent in the second.

The four peer-mean assignments preserve the negative sign, with ranges of −2.369 to −1.373 and −1.083 to −0.409 points. In contrast, shifting every missing-link common component by observed within-cell residual quantiles produces much wider bounds, −9.896 to 7.106 and −11.583 to 9.303 points. These are residual-shift bounds, not alternative peer-mean assignments.

The imputation diagnostics counsel against treating these totals as headline estimates. In masked-price validation, unweighted prediction correlations are only 0.028 and 0.030; value-weighted correlations are 0.129 and 0.271. Mean value-weighted prediction errors are 0.032 and −0.006 log units. Because the fixed basket has 15.5 and 17.7 percentage points more imputed weight than the actual basket, a common imputation bias of only 6–9 log points can reverse a gap. Moreover, the masked observations are surviving links, while the prices that require imputation belong to links that do not survive. Excluding each quarter's largest 0.1 percent of relationships reduces the gaps to −0.446 and −0.374 points; excluding the largest 1 percent reduces them to −0.168 and −0.127 points. The 2019 placebo gap, −1.294 points, also exceeds the comparable 2021 gap of −1.145 points.

Turnover itself is economically large. Relationships absent from the baseline account for 32.6 and 39.0 percent of actual value; re-entry of links observed before the baseline accounts for 9.8 and 12.7 percentage points of those totals. Baseline relationships no longer active account for 31.7 and 38.8 percent of fixed value, including 10.4 and 14.6 points of temporary dormancy. New origins within continuing importer–product pairs account for 7.4 and 9.0 percent of actual value. The largest 1 percent of relationships carry about 59–62 percent of

expenditure in both networks.

Table E.1: Imputed full fixed-network sensitivity exercise

Episode	Counterfactual	Actual	Fixed	Actual − fixed	Actual observed	Fixed observed
2014–16	Imputed full fixed network	−7.839	−6.424	−1.415	59.9	44.4
2021–22	Imputed full fixed network	10.338	11.421	−1.083	66.4	48.7

Notes. Entries are episode averages of quarterly four-quarter dollar log changes, in percentage points; they are not cumulative inflation. The same actual-minus-fixed gaps apply in pesos. Both indexes price the union of the actual and baseline networks with the same imputed price-growth function. The final two columns report the shares receiving observed relationship prices; the remainder is imputed. These estimates are not directly comparable with the observed-support counterfactuals in Table 3, and the gaps must not be added.

The peer-mean estimates point in the same direction as the observed-price exercises, but their magnitude is not stable to plausible changes in unobserved prices or to excluding dominant relationships. Weak masked-price performance and sizable pre-episode placebo gaps reinforce that this exercise is assumption-dependent sensitivity evidence, not a causal effect of episode-specific network adjustment.