Did Foreigners Pay America's Tariffs? Quantity Discounts, Scale Economies and Incomplete Pass-Through*

Sharat Ganapati
Georgetown University & NBER

Colin J. Hottman
Federal Reserve Board

January 2026

Abstract

Transaction-level quantity discounts are a pervasive feature of US trade, shaping both price variation and tariff incidence. Using administrative microdata, we show that transaction-level scale economies drive the majority of price dispersion within product categories. These discounts reflect supply-side scale economies rather than market power. Accounting for these micro-level economies resolves a key puzzle: while observed import prices rose one-for-one with 2018 US tariffs, we show this was driven largely by the loss of scale economies as transaction sizes collapsed. Controlling for this scale effect, the strategic pass-through of tariffs to prices falls to 60 percent, implying foreign exporters absorbed a significant share of the burden through reduced markups.

*We thank Nuno Limao, Charly Porcher, Daniel Xu, Justin Pierce, Davin Chor, Esteban Rossi-Hansberg, Frederic Warzynski, and Xiang Ding for early conversations and Nahim Bin Zahur for discussion. Any views expressed are those of the authors and not those of the US Census Bureau or the Federal Reserve Board. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center. (CBDRB-FY25-P2193-R12042, R12792)

1 Introduction

Determining whether foreign producers or domestic buyers ultimately bear the burden of tariffs is a central question in international economics. As has been long understood, a critical element in determining incidence is the pass-through of tariffs into prices (Jenkin, 1872). In this paper, we revisit tariff incidence and pass-through in the 2018-2019 period, taking into account an underappreciated feature of US trade: transaction-level quantity discounts. Using comprehensive US microdata, we show that these discounts account for the majority of price variation within narrowly defined product categories, and are inconsistent with the typical interpretation of the 2018-2019 period. Our findings reveal that accounting for quantity discounts reduces measured tariff pass-through, with foreign exporters shouldering a much larger share of the tariff burden than previously recognized.

In our empirical analysis, we distinguish between factors that affect the level of prices (such as firm productivity or market power) and those that affect the scale of prices (how unit prices change with transaction size). If scale reflects price discrimination, tariffs may simply transfer rents. If scale reflects real resource costs (e.g., logistics), shrinking order sizes represent a rise in real unit costs - a loss of resources. While market power may shift price levels, we find that it does not explain the steepness of the quantity discount schedule. Specifically, we estimate that a 1% increase in transaction size leads to a 0.29% reduction in price—a "scale effect" that is robust across relationships and driven by genuine supply-side economies rather than price discrimination.

We aggregate the transaction-level data to aggregate unit values and run a series of tariff pass-through regressions. We find that tariff pass-through to post-tariff buyer prices, as estimated in prominent work, is about 100 percent. However, this result masks two important underlying forces. First, we find that these measures of pass-through (near 100 percent) conflate two opposing forces: a strategic price cut by exporters and a mechanical price rise due to the loss of scale economies. We estimate that the 'scale-free' pass-through—reflecting the shift in the exporter's price schedule—is only 60 percent. Second, had exporters not cut markups, our results imply that observed tariff pass-through would have exceeded 100 percent given the downward sloping curve. While we do not account for any foreign retaliation, this provides a micro-founded baseline for future research.[1]

A striking feature of both domestic and international markets is the substantial price variation observed across seemingly identical products—a fact well documented in the literature (e.g., Stigler, 1961; Halpern and Koren, 2007; Atalay, 2014; Fontaine, Martin, and

[1]Extending our results to China's retaliatory tariffs and conducting a general equilibrium analysis are beyond the scope of this paper.

Mejean, 2020; Bornstein and Peter, 2024; Burstein, Cravino, and Rojas, 2024). Bornstein and Peter (2024) and Meleshchuk (2017) interpret these price differences entirely through the lens of markup differences across orders of different transaction sizes. Similarly, Alviarez, Fioretti, Kikkawa, and Morlacco (2023) assume market power and Burstein et al. (2024) assume exogenous wedges. We find little empirical support for these mechanisms in US trade. Instead, our results align with the distribution and operations research literature (Munson and Rosenblatt (1998), Munson, Jackson et al. (2015), Hornok and Koren (2015)), and find that quantity discounts are driven by scale economies and fixed costs that must stem from the underlying production and demand processes[2]. Our analysis demonstrates that these discounts are not merely anecdotal but are a pervasive and quantitatively important feature of US trade.

The implications for tariff pass-through are profound. Estimates of the incidence of the 2018-2019 tariffs are that they fell almost entirely on US buyers at the border, with pass-through rates nearly unity (Amiti, Redding, and Weinstein, 2019; Gopinath, Itskhoki, and Rigobon, 2010; Cavallo, Gopinath, Neiman, and Tang, 2021; Fajgelbaum, Goldberg, Kennedy, and Khandelwal, 2020). To explain this result in a standard model setup requires that exporting firms have constant markups and that the export supply curve facing the US is perfectly elastic (resulting in a flat marginal costs curve). In that setup, exporters bear none of the tariff burden. In contrast, we estimate that the export supply curve facing the US slopes downward (as is needed to explain quantity discounts).

Our approach demonstrates that the way scale economies aggregate across transactions is fundamental to understanding overall tariff pass-through. By using detailed transaction-level data, we show that micro-level quantity discounts and scale economies, when aggregated, directly determine the observed market-level incidence of tariffs. This perspective moves beyond traditional models that assume uniform pricing or infer scale economies from aggregate data (e.g., Antweiler and Trefler (2002); Bartelme, Costinot, Donaldson, and Rodriguez-Clare (2019); Lashkaripour and Lugovskyy (2023)). We extend the standard pass-through framework (e.g., Weyl and Fabinger (2013); Amiti et al. (2019); Ganapati, Shapiro, and Walker (2020)) by showing that aggregate pass-through is not a primitive, but an outcome shaped by the distribution of transaction sizes and the prevalence of non-linear pricing. This approach clarifies how micro-level pricing behavior translates into macro-level effects.

Our single correction helps resolve three puzzles in international trade. First is the

[2]Our approach is complementary to models that emphasize transaction costs and search frictions (e.g., Allen (2014); Krolikowski and McCallum (2025)), but we do not explicitly model these frictions. Instead, we capture their effects to the extent that they manifest as scale economies or quantity discounts in observed prices.

pass-through puzzle: recent empirical work finds complete tariff pass-through to US import prices (Fajgelbaum et al., 2020; Cavallo et al., 2021), a result that in standard models implies foreign export supply is perfectly elastic. This contradicts the widespread belief that the U.S. is not a small open economy, and the views of the optimal tariff literature (Broda, Limao, and Weinstein, 2008; Ossa, 2014) that U.S. tariff pass-through should be incomplete. We show that supply is actually downward-sloping ($\gamma < 0$); pass-through only *appears* complete because plummeting transaction volumes push unit costs up the supply curve, masking the strategic price cuts by exporters.

Second is the exchange rate disconnect puzzle: while tariff prices appear to move one-for-one, exchange rate pass-through is notoriously incomplete, typically around 0.5 (Gopinath et al., 2010; Amiti, Itskhoki, and Konings, 2014). By correcting unit values for quantity distortions, we align tariff pass-through estimates (≈ 0.6) with these established exchange rate findings, suggesting a unified pricing behavior across shocks.

Third is the incidence puzzle: while US import prices suggest domestic consumers bore the entire tariff burden, independent evidence from foreign production sites reveals significant welfare losses for exporters (Chor and Li, 2024). This creates a tension: if exporters fully passed on the tariffs, why did their local economies suffer such disproportionate losses? We resolve this by showing that pass-through was incomplete. Exporters lowered their effective markups and prices through quantity adjustments, thereby absorbing a significant share of the incidence, consistent with the observed contraction in foreign activity.

The rest of the paper proceeds as follows. We first decompose transaction-level prices and show that the size of transaction explains a very large portion of the dispersion in observed prices. Second, we show this is a largely supply-side phenomenon, where shifts in demand identify the transaction-level supply elasticity. Third, we show that variation in quantity discounts is driven by observable cost factors (shipping costs and within-firm transactions), as well as being higher in markets with low levels of demand substitutability. Fourth, we consider aggregation and show that reduced-form estimations of economies of scale and pass-through require additional interpretation, relative to our structural baseline with quantity discounts.

2 Conceptual Framework

Analyzing tariffs with quantity discounts requires a framework accommodating flexible supply slopes and potentially non-linear pricing (Stole, 2007). Classic analysis typically assumes upward-sloping supply and uniform pricing (Weyl and Fabinger, 2013). We

review the uniform pricing benchmark before formalizing our transaction-level decomposition. This approach sets the stage for our later empirical analysis, which will assess the prevalence and consequences of quantity discounts in US import data.[3]

Under uniform pricing and specific taxes, pass-through ($\rho = \partial p/\partial t$) is well-defined.[4] Incidence is the ratio of downstream to producer welfare changes: $I \equiv (\partial DS/\partial t)/(\partial PS/\partial t)$. Under monopoly, $I = \rho$; under perfect competition, $I = \rho/(1-\rho)$. Conditional on market structure, pass-through is a sufficient statistic for local incidence.[5]

Non-uniform pricing complicates this logic. Assume a continuum of customers optimally choosing quantities.[6] If a monopolist implements first-degree price discrimination, they price along the demand curve, capturing all surplus. With a tariff, a single pass-through figure is undefined, as each customer type faces a different price and pass-through rate. Aggregate pass-through is uninformative. In this case, incidence is mechanically 0 ($I = 0$), as all surplus loss falls on the producer. Conversely, under perfect monopsony where the buyer extracts all surplus, a tariff extracts all buyer surplus, and incidence is infinite ($I \to \infty$).

We decompose the price of transaction t between seller i and buyer j for variety v and quantity q with tariffs τ as:

$$\log P_t = \log \tilde{p}_t + \log p(q_t) = \log \tilde{c}_t + \log \tilde{\mu}_t + \log c(q_t) + \log \mu(q_t) + \log(1+\tau_t). \quad (1)$$

Normalizing $c(1) = \mu(1) = 1$, level terms $\tilde{p} = \tilde{c}\tilde{\mu}(1+\tau)$ capture product- and firm-level heterogeneity, while scale terms $p(q) = c(q)\mu(q)$ reflect variation with size. This distinguishes scale economies from level effects. Since costs and markups vary with quantity, a single price is undefined. However, we can still analyze the local effects of a policy by aggregating up from the transaction level.[7]

Figure 1 shows real-world examples of quantity discounts from Alibaba, a large online marketplace and Amazon Business. In both panels, we see that as the quantity purchased increases, the unit price decreases significantly. In the first example (Figure 1a), we'd say

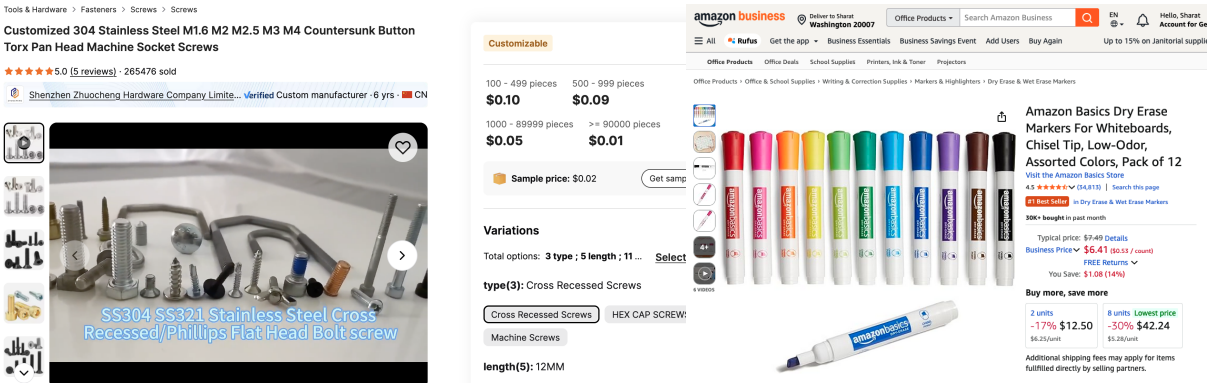
[3]While these results have been explored theoretically (Maskin and Riley, 1984), empirical study is largely neglected. (Exceptions include Meleshchuk (2017) in trade and Verboven (2002) in industrial organization).

[4]For exposition, we use a specific tax. An ad-valorem tax, such as a tariff, changes the derivations, but the same logic applies (Bishop, 1968; Cheung, 1998)

[5]The relationship between incidence and pass-through depends on the assumed market structure. While $I = \rho$ under monopoly and $I = \rho/(1-\rho)$ under perfect competition, intermediate cases of oligopoly can also be characterized. For example, under symmetric Cournot competition, incidence can be expressed as a function of both the pass-through rate and a "conduct parameter" θ that captures the degree of market power, as formalized in the general framework of Weyl and Fabinger (2013).

[6]We focus on quantity, but the framework is isomorphic to quality (Melitz, 2003).

[7]We assume scale terms are common across transactions; we test this in the empirical section.



(a) Screw Bolts on Alibaba (b) Markers on Amazon Business

Figure 1: Real World Examples of Quantity Discounts

that $\tilde{p} = 0.10$ and $p(q)$ decreases from 0.10 to 0.04 as quantity increases from 1 to 100. The elasticity of $p(q)$ to q is approximately 0.35.

2.1 Framework: Aggregate Scale Economies

Equation (1) defines transaction-level scale economies. We also consider how aggregate prices and costs respond to shocks. Using the scale-free price $\tilde{p}_t = \tilde{\mu}_t \tilde{c}_t$, we examine how an aggregate price index $\mathbf{P}$ changes with non-linear pricing.

We define the aggregate price index $\mathbf{P} = \sum_{t \in T} w_t \tilde{p}_t$, where $\tilde{p}_t = P_t / p(q_t)$ is the scale-free price and w_t are weights. This separates cost/markup changes from compositional effects, decomposing aggregate scale economies $\mathcal{S}_P$ into level effects ($\tilde{p}$) and scale effects ($p(q)$).

2.2 Framework: Tariff Incidence

We analyze the local incidence of tariffs on downstream buyer surplus, producer profits, and government revenue. Buyer surplus changes are derived by integrating under the demand curve (Weyl and Fabinger, 2013):

$$DS(\mathbf{p}) \equiv \int_{t \in ijvt} \int_{P_t}^{\infty} D(z) dz dt,$$

such that to a local approximation, defining a scale-free pass-through as $\tilde{\rho} \equiv \frac{d \log \tilde{p}_t}{d \log(\tau + 1)}$,

$$\frac{dDS}{d(\tau + 1)} = -\int_t \tilde{\rho} \frac{E_t}{1 + \tau} dt \tag{2}$$

where E_t represents expenditures.

This computation does not require demand curvature, which is only needed for global results and follows from individual buyers internalizing how changes in orders affect their prices.

Suppose we use unit values where $\rho \equiv \frac{d \log P_t}{d \log(1+\tau)}$ as a reduced-form pass-through. Then, pass-through estimated using unit values conflates strategic price changes with mechanical scale effects:

$$\tilde{\rho} = \rho - \frac{d \log p(q_t)}{d \log q_t} \frac{d \log q_t}{d \log(1+\tau)}.$$

If scale economies are large and quantities fall with tariffs, pass-through will be overestimated. Conversely, if we use ρ instead of the true pass-through of scale-free prices $\tilde{\rho}$, the change in downstream surplus will be overestimated.[8] However, the unit pass-through ρ still plays a role, as it is linked to producer surplus, which we now turn to.

Producer surplus is defined as:

$$PS \equiv \int_{t \in ijvt} [R(t) - C(t)] \, dt,$$

where $R(t)$ is the revenue from a transaction and $C(t)$ are the costs of a transaction.

We can then differentiate with respect to the tariff rate:

$$\frac{dPS}{d(\tau+1)} = \int_t \left[\frac{dR(t)}{d(\tau+1)} - \frac{dC(t)}{d(\tau+1)} \right] dt. \tag{3}$$

Revenue is observable; costs are not. Using Equation 1, we decompose revenue into costs and markups. We require: (1) scale economies and price discrimination $(c(q), \mu(q))$; (2) the response of revenue and quantity to tariffs; and (3) base costs $\tilde{c}$ or markups $\tilde{\mu}$. We estimate the first two empirically and calibrate the third from literature.

Revenue changes can be expressed in terms of local elasticities:[9]

$$\frac{dR}{d(1+\tau)} = \left[\frac{d \log \tilde{p}}{d \log(1+\tau)} + \left(1 + \frac{d \log p(q)}{d \log q}\right) \frac{d \log q}{d \log(1+\tau)} - 1 \right] \frac{R}{1+\tau}.$$

[8] Appendix A.1.1 allows for $p(q_t)$ to directly vary with $1+\tau$.

[9] We define prices as the tariff-inclusive price (FOB price plus tariffs). We revisit the difference with the freight inclusive price (CIF) in Section 5.4. Additionally, the appendix shows that if we know market conduct, the incidence ratio $I = \frac{dCS/d \log T}{dPS/d \log T}$ is: $I = \frac{-\text{Expenditure} \times \tilde{\rho}}{-\text{Revenue} \times [1-(1-\theta)\rho]} \approx \frac{\tilde{\rho}}{1-(1-\theta)\rho}$.

Cost changes are similarly decomposed:

$$\frac{dC}{d\left(1+\tau\right)} = \left[\frac{d\log\tilde{c}}{d\log\left(1+\tau\right)} + \left(1+\frac{d\log c\left(q\right)}{d\log q}\right)\frac{d\log q}{d\log\left(1+\tau\right)}\right]\frac{C}{1+\tau}$$

We decompose the price-quantity elasticity into scale economies $\gamma_c \equiv \frac{d\log c(q)}{d\log q}$ and price discrimination $\gamma_\mu \equiv \frac{d\log \mu(q)}{d\log q}$.

We use the 2018-2019 US tariffs to identify changes in scale-free prices, quantities, and costs. Recovered scale economies allow us to decompose price changes and compute incidence and surplus changes.

Government revenue is: $G \equiv \int_t \tau R(t)$. The change in government revenue with respect to the tariff rate is:

$$\frac{dG}{d\left(1+\tau\right)} = \int_t \left[\tau\frac{dR\left(t\right)}{d\left(1+\tau\right)} + R(t)\right].$$

Global incidence requires estimating supply and demand curves and integrating, necessitating functional form assumptions (see Ganapati and Hottman (2025)).

Before we can implement the framework above, we first turn to a basic question. What is the nature of price variation in the underlying data, and as decomposed by equation (1)? Is there even a correlational link between transaction size and price?

3 Decomposing Prices

3.1 Data

We focus on the universe of US imports from abroad. We consider all imports that exceed the *de minimis* threshold and are reported to the US Customs and later processed by the US Census to add firm identifiers. This combined database, the Longitudinal Firm Trade Transactions Database (LFTTD), is the source for every shipment transaction entering the US via air, water, or land.

The data consider exporters (labeled as the "Manufacturer ID" from the address on the customs record, but may be a wholesaler or another intermediary) shipping products to downstream US firms (Kamal and Monarch, 2018). These downstream US firms are a mix of manufacturers, wholesalers, and retailers. Our units of observation are individual transactions of an exporting firm, importing firm, 10-digit US Harmonized Tariff Schedule code (HTS 10-digit), port of entry, and date of entry. For each transaction, we observe the

value, shipping and freight charges, method of transport, tariffs and customs charges, a standardized measure of quantity within that HTS code, and the weight of the shipment. Of note, quantities and weights are often equivalent for many HTS codes, as goods are often sold by weight.

Prices are computed as value/quantity, though further analysis will also compute prices inclusive of tariff and freight charges. We decompose prices for 2017 and report robustness using 2012 data. In 2017, 50% of trade transactions occurred between parties that trade more than 100 times within an HTS 10-digit code in a year.

Alternative Data For domestic trade, we consider the US Census and Department of Transportation Commodity Flow Survey (CFS), focusing on 2012. See Appendix C.1 for details. The downside of this data is the lack of information on the identity of the buyer. Collectively these two datasets cover the entire universe of final and intermediate physical goods consumed in the United States.

3.2 Decomposing Import Prices

We demean log variables by variety (HTS-10).[10] We decompose prices into a scale component (γ) and a level component (μ). The relationship for transaction t is:

$$\log p_{i\to j,t,v} = \gamma \log q_{i\to j,t,v} + \mu_{i\to j,v} + \xi_{i\to j,t,v}, \tag{4}$$

where γ is the quantity scale elasticity and $\mu_{i\to j,v}$ is the relationship-specific price level. μ absorbs time-invariant factors (markups, quality, bargaining), isolating the scale effect γ. Index v is HTS 10-digit and source country; i is seller; j is buyer.[11]

Table 1 decomposes 2017 import prices for 60 million residualized transactions. A uniform quantity relationship explains 44.6% of price variation (Spec 1). Allowing quantity discounts to vary by HTS 6-digit (following the logic of Rauch (1999)) code explains 57.6% of variation (Spec 2).[12]

A stricter test looks within Buyer-Seller-Variety relationships. Seller-Variety fixed effects account for 73% of the price variation in the data. Buyer-Variety fixed effects account for 50% of the price variation. These highly disaggregated controls absorb 73.3% of the variation in prices (Specification 3). Adding a uniform quantity discount explains 38.3%

[10] We explore time variation, demeaning by HTS 10-digit and month. Results are largely consistent.

[11] Samples are aligned to include only observations present in all decompositions, removing singletons (dropping about 5% of observations). Additionally, demeaning either before or after the logarithmic transformation yields qualitatively similar results.

[12] 10-digit and 6-digit level HTS level quantity discounts are similar.

of the variation (specification 4; 17.4% directly and 20.9% through the covariance). Varying discounts by HTS 6-digit strengthens this (Spec 5). Even with millions of relationship fixed effects, transaction size explains a large share of price variation.

We compare this to yearly aggregate measures (market power, total annual volume) in Specifications 6-10.

The aggregate quantity (within variety), bought by a particular buyer, sold by a particular seller, and sold from a particular seller to buyer, is defined:

$$\log\left(q_{i,v}\right) \equiv \log \sum_{i',v'=i,v} q_{i'\to j,t,v'}, \qquad \log\left(q_{j,v}\right) \equiv \log \sum_{j',v'=j,v} q_{i\to j',t,v'},$$
$$\log\left(q_{i,j,v}\right) \equiv \log \sum_{i',j',v'=i,j,v} q_{i'\to j',t,v'}. \tag{5}$$

Aggregate volume measures explain only 6.5% of variation (Spec 6), compared to 57.6% for transaction-level quantity (Spec 2). Adding aggregate measures to transaction-level quantity adds minimal explanatory power, increasing the R-squared from 44.6% (Spec 1) to just 45.0% (Spec 7). This suggests that discounts are driven by per-shipment scale economies — such as in logistics, production, or handling — rather than by bargaining power derived from aggregate purchasing volume.

We measure traditional market power using bilateral market shares (Bernard, Dhyne, Magerman, Manova, and Moxnes, 2022):

$$s_{i\to j,v} = \frac{q_{i,j,v}}{q_{i,v}}, \qquad s_{i\leftarrow j,v} = \frac{q_{i,j,v}}{q_{j,v}}. \tag{6}$$

Shares account for only 0.8% of variation (Spec 8). Specification 9 confirms minimal additional explanatory power. Compared to seller/buyer fixed effects (Spec 10, 11), relation-level aggregates have minimal power. [13]

Are these effects driven by multiple transactions and shipments reflecting larger orders, where a large importer schedules multiple shipments at once? We do not observe orders, simply transactions. To explore this, we aggregate up to a potential order level, defined as all transactions between a buyer and seller within a variety over a given window of time. Specification 7 of Table 1 shows that using total yearly purchases has almost no extra explanatory power than simply using transaction-level quantities. Weaker results hold for monthly and quarterly aggregation.

[13] In further robustness checks, we show that the price variation associated with other relationship-level outcomes, such as relationship length and intra-firm trade, has even lower explanatory power.

Table 1: Correlating Price Variation: Decomposition of Log(Price) Variation

	Specification		Variance Decomposition			
	Controls	Fixed Effects	Controls	Fixed Effects	Covariance	Residual
1	$\log q$		44.6%			55.4%
2	$\log q$ x HTS6		57.6%			42.4%
3		Seller-Buyer-Variety		73.3%		26.7%
4	$\log q$	Seller-Buyer-Variety	17.4%	41.6%	20.9%	20.1%
5	$\log q$ x HTS6	Seller-Buyer-Variety	31.5%	32.0%	20.2%	16.3%
6	Aggregate Quantities		6.5%			93.5%
7	+ $\log q$		45.0%			55.0%
8	Relationship Shares		0.8%			99.2%
9	+ $\log q$		44.6%			55.4%
10		Seller-Variety, Buyer-Variety		72.5%		27.5%
11	Relationship Shares	Seller-Variety, Buyer-Variety	1.5%	72.6%	-1.5%	27.4%

Notes: This table decomposes 2017 import transaction-level price variation after demeaning log prices and quantities by country-origin and variety (HTS 10-digit) fixed effects. Sellers are designated at the Manufacturer ID and variety. Buyers are designated at the domestic firm ID and variety. For consistency, the sample is fixed to remain constant over all specifications. See the text for full details and specification.

Alternative Sample: Domestic Trade Data In domestic trade data we broadly find similar trends to the international trade import data. Appendix Table 24 replicates the exercise. Within tightly defined varieties, a simple log-linear quantity discount explains between 30-40% of all price variation. See Appendix C.1.1 for further details.

These descriptive regressions show transaction size predicts price, even with fixed effects. Scale effects correlate to over half of price variation, distinct from level effects. We next establish causality, isolating supply-side scale economies from demand effects.

4 Scale Economies and Quantity Discounts in Purchases

Is the price-quantity link driven by supply-side scale economies or downward-sloping demand? We isolate the supply-side to recover the buyer's price schedule. Identifying the transaction-level supply slope is challenging. Negative price-quantity correlation could reflect scale economies or a demand curve.

While demand estimation is a staple of applied economics, credibly identifying the slope of the supply curve—particularly at the transaction level—presents unique challenges. A negative correlation between price and quantity could reflect scale economies (a downward-sloping supply curve) or simply the tracing out of a demand curve.

We adopt a log-linear specification, a choice supported by robustness checks that reject alternative functional forms and also naturally fitting the local incidence framework. We decompose the price of a transaction t between seller i and buyer j for variety v and

quantity q as follows:

$$p_{i\to j,t,v} = \tilde{p}_{i\to j,v} q_{i\to j,t,v}^{\gamma} \xi_{i\to j,t,v}. \tag{7}$$

$\tilde{p} = \tilde{\mu}\tilde{c}$ is the unit price shifter. ξ captures unobserved quality/measurement error. $\gamma = \gamma^c + \gamma^\mu$ combines cost-scale (γ^c) and price-discrimination (γ^μ) elasticities.

The demand equation is:

$$q_{i\to j,t,v} = \tilde{q}_{i\to j,t,v} f\left(p_{i\to j,t,v}\right), \tag{8}$$

where $\tilde{q}$ is a demand shifter. Recovering γ requires variation in q uncorrelated with $\tilde{p}$. In our identification strategy, we adapt Arrow, Harris, and Marschak (1951) and Baumol (1952): firms with fixed ordering and inventory costs vary order sizes with demand realizations (Appendix A.8). Following Berry and Haile (2021), we use high-frequency demand shocks $\tilde{q}_{i\to j,-t,v}$ to avoid endogeneity with unobserved quality $\xi_{i\to j,t,v}$.

In Appendix B.3, we also put explicit structure on the demand following Feenstra (1994). In this situation with high frequency data, we assume that variation in residual supply and demand across different sellers within a buyer and across time (after double differencing) are orthogonal to each other. Broadly speaking both results find similar results for the curvature of the pricing curve.

We now detail our instrumental variable and fixed effect strategy.

4.1 Scale Pricing With Instrumental Variables and Fixed Effects

The estimating relationship is:

$$\log p_{i\to j,t,v} = \gamma \log q_{i\to j,t,v} + \mu_{i\to j,v} + \xi_{i\to j,t,v}. \tag{9}$$

We use buyer-seller-variety fixed effects, identifying off within-relationship quantity variation. We micro-found downstream demand shocks affecting the optimal inventory stocking of a downstream firm (Arrow et al., 1951), with firms ordering different levels of a variety over the course of a year. This model also produces variation in order frequencies, consistent with the lumpiness of international trade flows documented in inventory-based models (Alessandria, Kaboski, and Midrigan, 2010). The scale elasticity γ, governs how quickly prices change as quantity ordered increases within a buyer-seller relationship.

This strategy is in the spirit of Hillberry and Hummels (2013), but differs in a critical element - we consider transactions and abstract away from any measures of aggregates of transactions or direct measures of classic market power.

Causal identification faces two threats: unobserved relationship characteristics and simultaneous supply curve shifts. Unobserved quality variation (e.g., different iPhone models under one HTS code) could bias results. If quantity varies inversely with quality, this creates a mechanical negative relationship, opposite to the standard IO bias in demand estimation. This case resembles third-degree price discrimination or Alchian-Allen effects (Hummels and Skiba, 2004). We require a time-varying demand shock at the buyer level to address these threats.

Unobserved Characteristics Are we mismeasuring quality within even the highly disaggregated HTS-10 categories and a trading pair of downstream buyers and factories? For example, a buyer may purchase 10,000 iPhone 16s at $500 and 1,000 iPhone 17s at $1000 and both under HTS code "8517.14.0050". If goods are sold in efficiency units, we need a quantity shifter, uncorrelated with quality ξ.

This would be consistent with some degree of third-degree price discrimination. It also would be consistent with a form of the Alchian–Allen "shipping the good apples out" effect (Hummels and Skiba, 2004). However, either effect would be a threat to identifying a form of the supply curve. (This exists for even relatively homogeneous products, such as concrete and gasoline). Essentially we need a time-varying demand shock, either at the buying firm level, or even better at the buying-firm HTS10-digit level or buying firm-sourcing country level.[14]

Measurement error A related issue is measurement error in the independent variable, quantity.[15] Measurement error has two potential links. First is in q itself and second in its implicit link to p through the fact it is generated using q. The first measure can mechanically induce attenuation bias if q is observed with classic measurement error. The second also induces attenuation bias; it creates a mechanically negative correlation between the error in the derived price and the mismeasured quantity q, and it too results in attenuation bias. As with standard measurement error, an instrument solves the additional error implicit in p through mismeasured q. We now turn to an instrumentation strategy.[16]

Instrumental Strategy Instrument validity depends on $Var(q, \xi)$. Positive covariance (larger quantities of higher quality) biases OLS upward; negative covariance biases it

[14]We experimented with exchange rate variation, but it had poor explanatory power at the weekly or monthly level, likely due to invoicing currency and lead time frictions.

[15]We assume that transaction values are reported without error, as tariffs are charged on them and are closely monitored.

[16]Our unit of observation is a transaction, thus relationships with more transactions are more influential, echoing the weight correction term in estimators such as Broda and Weinstein (2006).

Table 2: Recovering Quantity Discounts: OLS Results

	(1)	(2)	(3)	(4)	(5)
			$\log p$		
$\log q$	-0.83	-0.533	-0.469	-0.296	-0.287
	(0.0127)	(0.00249)	(0.00275)	(0.00328)	(0.00354)
R^2	0.766	0.913	0.446	0.797	0.81
Within R^2	0.766	0.531	0.446	0.253	0.244
Fixed Effects					
Variety		✓			
Buyer-Seller-Variety				✓	✓
Seller-Month-Variety					✓
Demeaned					
Month-Variety			✓	✓	✓
Country-Variety			✓	✓	✓

Notes: Round parentheses represent standard deviations clustered at the relationship level. Demeaning refer to process of regularizing all variables by country-origin and product variety (HTS 10-digit) fixed effects, as well as month and variety fixed effects. Sellers are designated at the Manufacturer ID and variety level. Buyers are designated at the domestic firm and variety level. See the text for full details and specification. Standard errors are clustered at the relationship.

downward. We require a high-frequency instrument. Assuming goods within an HTS-10/buyer pair are complements or substitutes, changes in other purchases act as instruments.

With our high degree of fixed effects, we need a fine-grained instrument for $q_{i\to j,t,v}$ that varies at the monthly or weekly transaction level. This instrumental strategy relies on the demand-side variation, exogenous shifts in buyer demand that trace out the supply curve.

We instrument transaction size $q_{i\to j,t,v}$ with the total quantity of variety v purchased by buyer j from other transactions in the same month:

$$IV_{i\to j,t,v} \equiv \log \sum_{t\neq t',t'\in(month_t)} q_{i\to j,t',v}. \tag{10}$$

This leverages buyer-level demand shocks, assumed orthogonal to seller-specific quality shocks.

The validity of this instrument rests on the assumption that buyer-level demand shocks for a variety v are correlated across purchase, while being orthogonal to seller-specific supply or quality shocks $\xi_{i\to j,t,v}$. For example, a downstream demand shock for a buyer's

final product would increase its demand for input v from all its suppliers, but would be uncorrelated with a temporary quality issue from a single supplier i. Because the instrument proxies for the buyer's total downstream activity, it shifts demand for inputs. Whether inputs are substitutes or complements determines the shape of the demand curve, but the shift (the instrument) remains correlated with quantity demanded from all suppliers, orthogonal to supplier-specific supply shocks. We provide a formal microfoundation for this instrument based on a buyer's problem in Appendix A.8. To control for time-varying supply shocks, we then include one further set of fixed effects.

Supply Shocks Is identifying variation coming off a shift in the supply curve, rather than a shift in demand? To ensure we identify the supply curve, we control for supply shocks as many trading relationships are long-lived: $\gamma_{i,v,t\in M_t}$, where M_t is the month of transaction t. These stringent fixed effects reduce sample size, dropping 30% of value and transactions.[17]

We also conduct robustness checks by varying the construction of the instrument, such as using alternative time windows (weekly, quarterly), looking only within a relationship, or excluding all purchases from the same supplier, with consistent results.

4.2 Transaction-level Supply: Import Regression Results

We present regression results for the scale elasticity γ (Equation (9)). Table 2 shows OLS results. The base regression shows -0.830 (Col 1), but clearly mixes supply and demand. Simple demeaning yields -0.469 (Col 3). Adding buyer-seller-variety fixed effects reduces this to -0.296 (Col 4). Including seller-month-variety fixed effects yields -0.287 (Col 5). Thus, a 10% transaction size increase lowers price by 2.9%.

Table 3 presents IV results using the instrument from Equation (10). The first stage is strong (Col 2). The second-stage estimate is -0.284, close to OLS (Col 3). A related instrument (constructed using only the total other purchases from the same manufacturer) yields similar results (Col 4), confirming $\gamma \approx -0.29$. Using both instruments finds similar results (Col 5). A final instrument that looks at purchase variation between manufacturers, that is leaving out all transactions within the focal relationship (Col 6), finds a slightly lower $\gamma = -0.200$, but with a weaker first-stage and wider standard errors.[18]

Table 4 summarizes heterogeneity across HTS 6-digit categories. We re-estimate our

[17] As robustness exercises, we controlled for time-varying country-level characteristics ($\gamma_{c,v,t}$) and relationship length within a year. Neither significantly affected the estimated coefficients. Using only arm's-length relationships yields similar results.

[18] We also experimented with weekly and quarterly versions of the instrument, finding similar results.

Table 3: Recovering Quantity Discounts: IV Results

	(1)	(2)	(3)	(4)	(5)	(6)
	log p	log q	log p			
log q	-0.268		-0.284	-0.315	-0.288	-0.200
	(0.00463)		(0.00151)	(0.00239)	(0.00152)	(0.0460)
Baseline Instrument		-0.773				
		(0.00446)				
R^2	0.800	0.642				
Within R^2	0.226	0.0527	0.226	0.219	0.225	0.218
First Stage F-Stat			30040	8425	10020	115
IV		First Stage	Baseline	Within-Relationship	3+4	Other Suppliers
Fixed Effects						
Buyer-Seller-Variety	✓	✓	✓	✓	✓	✓
Seller-Month-Variety	✓	✓	✓	✓	✓	
Demeaned						
Month-Variety	✓	✓	✓	✓	✓	✓
Country-Variety	✓	✓	✓	✓	✓	✓

Notes: Round parentheses represent standard deviations clustered at the relationship level. The sample for column (1) replicates the OLS specification for this sample. Demeaning refer to the process of regularizing all variables by country-origin and variety (HTS 10-digit) fixed effects, as well as month and variety fixed effects. Sellers are designated at the Manufacturer ID and variety level. Buyers are designated at the domestic firm and variety level. The J-statistic for overidentification is 5.722. See the text for full details and specification. Standard errors are clustered at the relationship.

preferred IV specification separately for each category to recover γ_v. In both the full OLS (Col 3) and IV specification (Col 5), most products exhibit significant discounts, with mean (-0.275, -0.284) and median (-0.223, -0.239) elasticities consistent with the aggregate. We use Column (5) as a baseline in our incidence calculations.

Domestic Data We estimate supply using domestic data from the CFS in Appendix C.2. Using a similar IV strategy, we find γ_v between -0.21 and -0.26.

Structural Estimation For robustness, we employ a structural approach to jointly estimate supply and demand elasticities, following Feenstra (1994). This method relies on the heteroskedasticity of orthogonal supply and demand shocks for identification rather than traditional instrumental variables. In Appendix B.3, we aggregate transactions to the monthly level for 6-digit HS categories and assume constant elasticities of substitution and supply. The structural estimates yield a median supply elasticity (ω) of approximately -0.23 to -0.37, which closely matches our reduced-form estimate of $\gamma \approx -0.29$. Additionally, we recover a median demand elasticity (σ) of approximately 3.4. The consistency between these structural estimates and our baseline results reinforces the valid-

Table 4: Heterogeneity in Quantity Discounts: Variation across Products

		(1)	(2)	(3)	(4)	(5)
				$\log p$		
$\log q$	median	-0.344	-0.267	-0.222	-0.241	-0.239
$\times$ Variety	mean	-0.367	-0.305	-0.275	-0.310	-0.284
	variance	0.211	0.197	0.217	1.160	0.245
R^2	mean	0.576	0.835	0.85		
First Stage F	mean				5.05	45.3
Fixed Effects						
Buyer-Seller-Product			✓	✓	✓	✓
Seller-Month-Product				✓		✓
Demeaned						
Month-Product		✓	✓	✓	✓	✓

Notes: This table summarizes the results from various specifications estimating the price elasticity with respect to quantity. Each column corresponds to a different model specification as detailed by the fixed effects and demeaning procedures applied.

ity of our supply-side identification. We note that while the structural approach provides valuable cross-validation, the assumed CES demand and monopolistic competition market structure are inconsistent with our later findings on pass-through. We refer readers to Ganapati and Hottman (2025) for a consistent treatment.

Validation of Approaches In Appendix B.1, we validate our baseline estimates by comparing the variety-level scale elasticities (γ_v) recovered from our OLS and IV strategies against those obtained from alternative methods. We find a strong positive correlation between our preferred estimates and those derived from a structural model of supply and demand (following Feenstra (1994)), as well as estimates based purely on observable shipping costs. The fact that shipping costs—a direct measure of logistics technology—exhibit scale economies that are highly correlated with our overall price elasticities reinforces our interpretation that the observed quantity discounts are driven by real resource costs rather than demand-side factors or model-specific assumptions.

Discussion: Transaction Level Supply Quantity discounts are pervasive. The negative price-quantity relationship is robust to OLS, supply shock controls, and IV strategies. We estimate the median $\gamma \approx -0.239$ across HTS codes. This aligns with Bornstein and Peter (2022) and Meleshchuk (2017), though interpretation (scale vs. markups) could differ.

5 The Mechanism Behind Quantity Discounts

Having established that the supply curve facing US buyers is downward sloping ($\gamma < 0$), we now determine its source. If γ reflects price discrimination, tariffs may simply transfer rents. If γ reflects real resource costs (e.g., logistics), shrinking order sizes represent an efficiency loss.

We distinguish between effects altering the *level* of prices (scale-free price $\tilde{p}$) and those affecting the *scale* elasticity (γ). We find the scale component is cost-driven. Relationship-level market power and bargaining explain little of variation in γ. Instead, evidence points to common, cost-based scale economies—arising from fixed transaction or shipping costs. This is supported by three facts: (1) market power proxies fail to explain heterogeneity in discounts; (2) intra-firm and arm's-length transactions exhibit identical schedules; and (3) cross-market variation correlates with cost and product attribute proxies rather than competition measures.

Recall the decomposition from Equation (1):

$$\log p(q) = \underbrace{\log \tilde{c} + \log \tilde{\mu}}_{\text{level}} + \underbrace{\log c(q) + \log \mu(q)}_{\text{scale}}.$$

Both the level terms ($\tilde{c}$, $\tilde{\mu}$) and the scale terms ($c(q)$, $\mu(q)$) can reflect a mix of cost and markup elements. The level terms capture product- and relationship-specific heterogeneity in marginal costs and markups, while the scale terms capture how costs and markups vary with transaction size.

We decompose the transaction-level elasticity $\gamma_{i \to j,v}$ into cost (γ^c) and markup (γ^μ) components:

$$\gamma_{i \to j,v} = \gamma^c_{i \to j,v} + \gamma^\mu_{i \to j,v}. \tag{11}$$

Following Stole (2007), optimal non-linear pricing links the markup component of the scale elasticity to the residual demand elasticity $\epsilon_{D,v}$ faced by the firm ($i \to j$):

$$\gamma_v = \gamma^c_v + (1 - 1/\epsilon_{D,v}). \tag{12}$$

If discounts reflect price discrimination, $\gamma_{i \to j,v}$ should vary with relationship characteristics affecting demand elasticity. Absence of such variation suggests the cost component γ^c_v dominates.

Empirical Strategy Our empirical strategy proceeds in three steps to disentangle cost-based scale economies from markup variation.

Step 1: Common quantity discounts. We begin by imposing that all seller-buyer pairs within a variety share the same cost-based scale elasticity: $\gamma^c_{i\to j,v} = \gamma^c_v$. Under this assumption, any observed heterogeneity in $\gamma_{i\to j,v}$ across relationships must reflect differences in markup behavior $\gamma^\mu_{i\to j,v}$. We test whether variation in market power—measured by seller concentration, buyer concentration, or bilateral market shares—systematically predicts this heterogeneity. If markup variation were the primary driver, we would expect relationships characterized by greater market power to exhibit systematically different quantity discounts. Alternatively, a common γ_v across relationships could also reflect a common markup component γ^μ_v arising from symmetric market frictions or search costs faced by all buyers within a market—an interpretation we explore further in Step 3.

Step 2: Offsetting effects. If we find little explanatory power from market power measures in Step 1, a mechanical possibility remains: firms with greater market power could exhibit both steeper markup schedules with respect to quantity (higher $\gamma^\mu_{i\to j,v}$) *and* flatter cost curves (lower $|\gamma^c_v|$), such that the two effects offset and leave the total scale elasticity $\gamma_{i\to j,v}$ relatively constant. To address this, we exploit variation in vertical integration. Intra-firm transactions—where transfer prices are often set to reflect cost—provide a natural benchmark for the cost component γ^c_v. If quantity discounts are similar for within-firm and arm's-length transactions after controlling for relationship fixed effects, this suggests that the common cost component γ^c_v dominates, and that variation in markup elasticity $\gamma^\mu_{i\to j,v}$ plays a minimal role.

Step 3: Search and transaction costs. Finally, we examine whether the common component γ^c_v itself reflects genuine production-side scale economies or instead arises from search and transaction costs on the buyer side. We test whether γ_v varies systematically with product-level characteristics that proxy for these costs. Specifically, we regress the estimated γ_v on measures of product substitutability (e.g., demand elasticity σ_v), market thickness (number of buyers and sellers), and transaction frequency. If quantity discounts were driven primarily by buyer-side search costs, we would expect larger discounts in markets with fewer suppliers, less substitutable products, or higher search frictions. Conversely, if discounts reflect supply-side scale economies—such as fixed costs per transaction or increasing returns in production—they should be relatively invariant to these demand-side characteristics and instead correlate with cost-related factors such as shipping scale economies or production technology.

Taken together, these three steps allow us to assess the relative importance of cost-based scale economies, markup variation due to market power, and search or transaction costs in explaining observed quantity discounts. Our findings, previewed in Tables 5 and 6, show that market power measures have minimal explanatory power, that intra-

firm and arm's-length transactions exhibit similar discounts, and that γ_v correlates most strongly with cost-side factors and product substitutability—consistent with the interpretation that $\gamma_{i\to j,v} \approx \gamma_v^c$ and $\gamma_{i\to j,v}^{\mu} \approx 0$.

Table 5: Decomposition of Price Variation by Interaction Term

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variable	$\log p_{i\to j,t,v}$								
Independent Variable	$\log q\times$ Related	$\log q\times \log(q_{ij})$	$\log q\times \log(q_i)$	$\log q\times \log(q_j)$	$\log q\times$ Output Share	$\log q\times$ Input Share	$\log q\times$ Rel. Length	$\log q\times$ HHI Seller	$\log q\times$ HHI Buyer
Panel A: Homogeneous Interactions									
Coefficient	-0.010	0.000522	0.000449	0.000418	0.00163	-0.0104	0.000609	-0.0169	-0.00751
	(0.00420)	(0.000195)	(0.000187)	(0.000165)	(0.00222)	(0.00402)	(0.000452)	(0.00779)	(0.00634)
Within R^2	0.0001	0.000197	0.000153	0.000157	0.0000106	0.000184	0.0000313	0.000175	0.000045
Panel B: Heterogeneous Interactions									
Mean Coef.	0.004	0.028	0.036	0.036	0.035	-0.0416	0.005	-0.202	-0.100
Median Coef.	-0.016	0.0195	0.021	0.021	-0.0226	-0.051	0.000	-0.072	-0.048
Variance Coef.	0.346	0.0833	0.0896	0.0896	1.05	0.34	0.129	0.866	1.15
99% Interval (lower)	-0.441	-0.0397	-0.0439	-0.0439	-0.707	-0.363	-0.00864	-1.01	-0.635
99% Interval (upper)	0.346	0.0833	0.0896	0.0896	1.05	0.34	0.129	0.866	1.15
Controls									
Variety Specific Quantity Discounts	✓	✓	✓	✓	✓	✓	✓	✓	✓
Fixed Effects									
Buyer-Seller-Variety	✓	✓	✓	✓	✓	✓	✓	✓	✓
Seller-Month-Variety	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: This table reports regressions of transaction-level prices on log quantity interacted with various relationship characteristics, while controlling for buyer-seller-variety fixed effects and seller-month-variety fixed effects. Each column represents a separate regression where the dependent variable is the interaction of log quantity with a specific characteristic. Standard errors are in parentheses.

5.1 Common Cost and Markup Structures

We first test if heterogeneity in quantity discounts $\gamma_{i\to j,v}$ correlates with market power, before we consider if there is a common markup component γ_v^{μ}. We estimate:

$$\log p_{i\to j,t,v} = \log \tilde{p}_{i\to j,v} + \gamma_v^{common} \log q_t + \beta \Gamma_{i,j,v} \log q_t + \epsilon_{i\to j,t,v}, \tag{13}$$

where $\Gamma_{i,j,v}$ represents relationship-specific deviations proxied by market power measures.

Table 5 presents the results. Panel A reports pooled specifications in the OLS framework; Panel B estimates separate regressions for each HTS 6-digit variety in the IV framework. Each column and row corresponds to a different market power measure, including vertical integration (related party status), bilateral market shares, relationship size (total quantity purchased), relationship length, and seller/buyer concentration (HHI weights, or squared market shares).

Panel A reports the results from the pooled specification. Across all specifications, market power measures have negligible explanatory power. For instance, the coefficients on 'Output Share' and 'Input Share' (Columns 5 and 6) are 0.0016 and -0.0104, respectively. Since shares range from 0 to 1, even a monopsonist would face a discount schedule nearly identical to a small buyer. The within-R^2 for these interactions is near zero.[19]

Panel B confirms this lack of heterogeneity across products. While there is heterogeneity across products (as seen in the variance of the coefficients), median interaction effects are consistently close to zero. For example, the median coefficient for the interaction with 'Output Share' is -0.023, and for 'Input Share' is -0.051. This lack of systematic variation suggests that heterogeneous markup behavior, $\gamma^{\mu}_{i \to j,v}$, is not a significant driver of quantity discounts. [20]

This supports the interpretation that the common component γ_v^{common} primarily reflects cost-based scale economies rather than heterogeneous price discrimination across relationships.

In Appendix B.2, we go a step further and recover a unique scale elasticity for every buyer-seller pair, rather than just interacting a common elasticity with relationship characteristics. We then regress these relationship-specific elasticities on measures of market power, including bilateral volume and market shares. Consistent with the interaction results, we find no significant correlation, reinforcing the conclusion that quantity discounts are not driven by relationship-specific markup variations. While buyers and sellers with and without market power have similar scale elasticities, they may have different reasons; there may be offsetting markup and cost effects. We explore this next.

5.2 Offsetting Effects and Common Costs

Could firms with market power have steeper markup schedules (γ^{μ}) but flatter cost curves (γ^{c}), masking the effect? Under this scenario, our finding of minimal heterogeneity in $\gamma_{i \to j,v}$ would not necessarily imply that markup variation is small—it could simply reflect a fortuitous cancellation of opposing forces.

To address this, we exploit variation in vertical integration. Intra-firm transactions,

[19] We avoid using a mix of revenue-based and quantity-based shares as regressors. For instance, one might consider using a seller's revenue share ($pq_{\text{relationship}} / \sum pq_{\text{seller}}$) and a buyer's quantity share ($q_{\text{relationship}} / \sum q_{\text{buyer}}$) to explain price variation. This approach is problematic because the dependent variable, log price, is mechanically defined as log revenue minus log quantity. Using a revenue-based regressor and a quantity-based regressor can create a spurious correlation.

[20] Relationship size and length (Antras and Helpman, 2004) show similar patterns: larger and longer-lived relationships exhibit marginally smaller quantity discounts, but these effects remain economically modest and add minimal explanatory power beyond the baseline average quantity discount (Col 7).

where transfer prices are set between divisions of the same firm and often reflect internal costs, provide a natural benchmark for isolating the cost component γ_v^c. If quantity discounts are similar for within-firm and arm's-length transactions after controlling for relationship fixed effects, this suggests that the common cost component γ_v^c dominates, and that variation in markup elasticity $\gamma_{i\to j,v}^{\mu}$ plays a minimal role.

Using the US Customs definition of related parties as our indicator of vertical integration, we test this hypothesis by estimating Equation (13) with an interaction term for related-party status.

Table 5, regression (1), shows that related-party status is associated with a statistically significant but economically negligible difference (1 percentage point) in the quantity discount. Intra-firm and arm's-length transactions face nearly identical schedules. Assuming transfer prices reflect costs, this stability supports the interpretation that $\gamma_{i\to j,v}$ primarily captures a common cost structure γ_v^c.

To further test the assumption that γ_v^c is common across all sellers, we relax the buyer-seller-variety fixed effects in regression (1) and instead include only buyer-variety and seller-variety fixed effects. This allows us to assess whether the same purchaser faces the same quantity discount from both related and non-related party sellers. Our results show that buyers indeed face similar quantity discounts regardless of whether the seller is a related party, which is consistent with the idea that quantity discounts reflect a common cost structure across transactions, not differential exercises of market power.[21]

5.3 Market-Level Variations and Search Costs

We next examine *cross-market* patterns by regressing the estimated quantity discount γ_v on market characteristics.

If discounts were driven by market-wide markup variation (γ_v^{μ}), we would expect them to be larger in less competitive markets. To test this, we regress the estimated quan-

[21] It is important to note that the particular laws governing the valuation of intra-firm transactions may complicate the interpretation of the related-party flag. Transaction values for related-party trade can be determined by two broad methods: one using a method of appraisal (e.g., actual cost plus a fixed profit share computed at the firm level), and the other requiring that pricing be consistent with industry pricing and comparable to pricing for non-related buyers. Under the first method, if prices are determined using actual costs plus a fixed markup, then the quantity discount for intra-firm trade would reflect true cost differences, absorbed by the fixed effects. Under the second method, if prices are determined using a reference system that implicitly incorporates a constant measure of market power, then differences between intra-firm and arm's-length pricing may not reflect true differences in market power due to legal constraints on reporting.

tity discount γ_v on various market-level characteristics Y_v:

$$\gamma_v^{common} = \alpha Y_v + \epsilon_v, \tag{14}$$

where γ_v^{common} is the recovered quantity discount variety v which indexes a particular variety v.

Table 6 considers measures of market concentration (HHI) and market thickness (number of buyers/sellers or buyer-seller pairs) and shows that concentrated markets have shallower quantity discounts. This is inconsistent with a market power story: if discounts were a tool for price discrimination, we would expect them to be steeper where firms have more power.

In contrast, cost and substitutability proxies are informative. Markets with more elastic demand (higher σ) exhibit little to no change in quantity discounts in our preferred IV specification (Panel B, Column 8). This null result implies that the demand-side environment has no bearing on the discount schedule, contradicting markup-driven explanations. Instead, the discount schedule appears invariant to demand conditions, consistent with a universal cost-based mechanism where standardized goods share similar logistic scale economies regardless of substitutability. Column (9) shows that markets with greater price dispersion—a proxy for search or transaction frictions—exhibit larger discounts.

This finding points towards an alternative mechanism rooted in search and transaction costs. When buyers face fixed costs to find and vet suppliers, or when sellers face fixed costs per transaction, quantity discounts naturally arise as a way to amortize these costs over larger orders. This creates a form of scale economy, either on the buyer's side (search) or the seller's side (transaction). Such a mechanism is consistent with our finding that discounts are larger for less substitutable products, where search costs are likely to be higher. However, a pure search cost story has trouble explaining the uniformity of discounts *within* relationships. If search costs are paid upfront to establish a relationship, we would expect the scale elasticity γ to be smaller for subsequent, within-relationship transactions. Our finding that γ is stable and significant even within long-term relationships suggests that per-transaction fixed costs (e.g., for logistics, invoicing, or quality verification) are a more likely driver than initial search costs. As further evidence, Column (9) shows that markets with greater price dispersion—a common proxy for higher search frictions (Marshall, 2020)—exhibit significantly larger quantity discounts, consistent with a model where transaction costs are a key determinant of the pricing schedule.

Table 6: Quantity Discounts and Across Market Variation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variable	Quantity Discount γ_v								
Independent Variable	HHI Seller	HHI Buyer	log(Buyers)	log(Sellers)	log(Pairs)	log(Trans.)	log(Value)	Demand log(σ)	Variance($\tilde{P}$)
Panel A: Dependent Variable: OLS Scale Parameter by HTS-6 Code; γ_v^{OLS}									
Coefficient	0.0892	0.0904	-0.015	-0.0101	-0.00904	0.00138	0.00156	-0.0141	-0.154
	(0.0453)	(0.0378)	(0.00638)	(0.00594)	(0.00594)	(0.00500)	(0.00508)	(0.00543)	(0.0128)
R-squared	0.002	0.002	0.004	0.002	0.002	0	0	0.002	0.086
Panel B: Dependent Variable: IV Scale Parameter by HTS-6 Code; γ_v^{IV}									
Coefficient	0.193	0.15	-0.0355	-0.0283	-0.025	-0.00867	-0.0104	0.000773	-0.199
	(0.0679)	(0.0485)	(0.00687)	(0.00618)	(0.00542)	(0.00713)	(0.00716)	(0.00903)	(0.0208)
R-squared	0.001	0.001	0.003	0.002	0.002	0.000	0.000	0.000	0.008

Notes: This table reports regressions of market-level characteristics on quantity discounts. The dependent variable in both panels is the estimated quantity discount parameter γ_v at the HTS-6 code level. Panel A uses the OLS estimate of γ_v, while Panel B uses the IV estimate. Each column regresses γ_v on a different market-level characteristic. Standard errors are in parentheses.

5.4 Shipping and Transportation Costs

A key advantage of the import data is that it separately reports shipping and freight charges, allowing us to directly test for scale economies in a major component of transaction costs. If the overall quantity discounts we observe are driven by cost-side factors, we should see similar patterns in these observable cost components.

Table 7 confirms this prediction. We find substantial scale economies in shipping: regressing the log of per-unit freight costs on the log of transaction quantity, within a buyer-seller-variety relationship, yields a scale elasticity of -0.383 (Column 5). This means a 10% increase in shipment size is associated with a 3.8% decrease in per-unit shipping costs. While declared shipping charges are only 3-5% of overall transaction values, this is qualitatively consistent evidence of physical scale economies. It supports an interpretation that the overall quantity discount schedule is driven by real cost factors rather than markup adjustments.[22] This aligns with evidence that the technology of international transport is characterized by significant increasing returns to scale (Hummels, Lugovskyy, and Skiba, 2009; Ganapati et al., 2024), providing a physical underpinning for the cost gradients we observe.

This elasticity is robust across specifications. Columns (4) and (5) control for the mode of transport (e.g., air, containerized sea, bulk carrier, land), isolating scale economies from differences in shipping technology. Our preferred specification in Column (5), which includes relationship, country-variety, and mode fixed effects, confirms a strong and sig-

[22]We leave the study of shipping scale to future work, tying in with both market power and economies of scale, as in Ignatenko (2020); Ganapati, Wong, and Ziv (2024).

nificant elasticity of -0.383. Column (6) allows this elasticity to vary by country-variety; the mean elasticity remains large at -0.355, with a standard deviation of 0.381 across varieties (in brackets), indicating that while there is heterogeneity, shipping scale economies are a pervasive feature. The results are also robust to alternative measures of transaction size. Using shipment weight in kilograms instead of standardized quantity units yields a nearly identical elasticity of -0.378 (Column 7). Using shipment value yields a smaller but still significant elasticity (Column 8). This consistent pattern across different measures and controls provides strong, direct evidence for cost-based scale economies in a key component of transaction costs.

Table 7: Decomposition of Shipping and Other Charges

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			$\log\left(\frac{\text{Charges}}{q}\right)$				$\log\left(\frac{\text{Charges}}{\text{kg}}\right)$	$\log\left(\frac{\text{Charges}}{\text{Value}}\right)$
$\log q$	-0.508	-0.403	-0.377	-0.481	-0.383	-0.355		
	(0.000126)	(0.000116)	(0.000155)	(0.000122)	(0.000109)	[0.381]		
$\log kg$							-0.378	
							(0.000133)	
log Shipment Value								-0.168
								(0.000159)
$\log q \times$			-0.0577					
Related Parties			(0.000232)					
R^2	0.766	0.932	0.934	0.854	0.95	0.784	0.683	0.489
Within R^2	0.219	0.185	0.188	0.215	0.196	*	0.128	0.0189
Fixed Effects								
Country-Variety	✓			✓	✓	✓	✓	✓
Relationship		✓	✓		✓			
Mode				✓	✓			

Notes: This table decomposes various measures of shipping charges. See text for details. Parentheses represent standard errors, and brackets represent other statistics.

5.5 Discussion

The evidence strongly supports a cost-based explanation. While observable shipping charges (freight) only account for a portion of the value, they crucially reveal that the *technological* elasticity of logistics is large (≈ -0.38). We infer that similar scale economies apply to unobservable order-level costs, such as batch production setups, warehousing handling, and administrative processing (Arrow et al., 1951). For example, if internal batch-level costs share the same physics as external freight, they would generate the large aggregate discounts we observe. Conversely, if the remaining discount were markup-driven, we would expect it to vary with market power, which it does not.

The scale elasticity is a market-wide feature, common to all transactions, rather than a

tool for price discrimination. We conclude that the firm-to-firm quantity discount elasticity is well-approximated by a market-level cost component: $\gamma_{i\to j,v} \approx \gamma^c_v$, with $\gamma^{\mu}_{i\to j,v} \approx 0$.

Furthermore, we will conduct robustness in which we assume that we are wrong and that all the variation in quantity discounts is due to markup variation. In some sense our finding limits the extent from which producer surplus can be extracted from buyers through quantity discounts and limits the exposure of producers to tariffs.

Our finding allows us to decompose observed prices into two distinct components: a cost-driven scale effect (γ) and a residual "scale-free" price level ($\tilde{p}$). In the next section, we analyze the behavior of this residual price level to distinguish between market power and aggregate scale economies.

6 Aggregate Implications

Since transaction-level quantity discounts are primarily cost-driven ($\gamma \approx \gamma^c$), we can isolate a "scale-free" price, $\tilde{p}_{i\to j,v}$, representing the price level net of transaction-size effects.

6.1 The Determinants of Scale-Free Prices

Relationship Scale Economies Do transaction-level economies aggregate to the firm or relationship level? Recent policy literature often estimates scale economies at the firm or higher levels (Lashkaripour and Lugovskyy, 2023; Bartelme et al., 2019; Farrokhi and Soderbery, 2020). Going beyond correlations, we test whether large buyers pay lower prices because of aggregate volume (firm scale) or simply because they make larger individual shipments (transaction scale).

We employ a first-difference strategy using 2011-2012 data to identify the relationship-level supply elasticity:

$$\Delta \log \tilde{p}_{i\to j,v} = \theta \Delta \log Q_{i\to j,v} + FE_{i,v} + \epsilon_{ij}, \tag{15}$$

where $\Delta \log Q_{i\to j,v}$ is the change in total quantity purchased by buyer j from seller i. We control for seller-variety fixed effects ($FE_{i,v}$) to account for supply-side shocks. We instrument for the change in quantity using a shift-share instrument based on aggregate demand shocks in downstream industries.[23]

Table 8 compares regressions using raw average prices ($\Delta \log p$) versus scale-free prices ($\Delta \log \tilde{p}$).

[23]Due to Census disclosure lags we use 2011-2012 data here instead of 2016-2017. Results are broadly similar.

Table 8: Relationship Scale Economies - Bilateral Price Changes

	(1) $\Delta \log p$	(2) $\Delta \log \tilde{p}$	(3) $\Delta \log p$	(4) $\Delta \log \tilde{p}$	(5) $\Delta \log p$	(6) $\Delta \log \tilde{p}$	(7) $\Delta \log p$	(8) $\Delta \log \tilde{p}$
$\Delta \log q$	-0.209	-0.0225	-0.222	-0.0229	-0.193	-0.0184	-0.115	-0.00114
	(0.005)	(0.001)	(0.006)	(0.001)	(0.006)	(0.001)	(0.04)	(0.002)
N	230000	230000	230000	230000	148000	148000	148000	148000
r^2	0.353	0.187	0.527	0.419	0.527	0.431	0.153	0.0005
Within r^2	0.197	0.006	0.21	0.006	0.182	0.004		
Fixed Effects								
Country-Variety	✓	✓						
Seller-Variety			✓	✓	✓	✓	✓	✓
First Stage F							1000	
Instruments							✓	✓
IV Sample					✓	✓	✓	✓

Notes: This table reports regressions of the change in aggregate quantities on the change in price. Odd columns use aggregate prices ($\Delta \log p$). Even columns use the scale-free price ($\Delta \log \tilde{p}$), which adjusts for transaction-level quantity discounts. The last two columns instrument for the change in quantity using the shift-share instrument described in the text.

When using raw prices (Columns 1, 3, 5, 7), we find a significant negative relationship between price and quantity, suggesting a level of aggregated scale economies. However, when we use the scale-free price $\tilde{p}$ (Columns 2, 4, 6, 8), this relationship largely disappears or is substantially attenuated. This implies that apparent aggregate scale economies are a mechanical reflection of transaction-level discounts. Large buyers pay less primarily because they transact in larger batches.

Our estimates align with Bartelme et al. (2019) and Duarte, Magnolfi, Quint, Sølvsten, and Sullivan (2025), who find more aggregated scale elasticities (0.1-0.17). Macroeconomic models assuming firm-level scale economies may mis-specify the mechanism if the true driver is at the transaction-level. We replicate these findings with domestic data in Appendix C.2.1. In Appendix B.5, we also decompose scale-free prices in the cross section.

Validating the Decomposition If our decomposition is correct, "level" effects like market power may reside in $\tilde{p}$, not γ. Table 9 confirms this. This is consistent with the literature on how there is wide variation in both how exporters react to various shocks (Berman, Martin, and Mayer, 2012) and in what determines price variation across relationships (Manova and Zhang, 2012; Kugler and Verhoogen, 2012; Basu and Fernald, 1997).

Panel A shows correlations between prices and measure of market power within Variety-Country. Panel B shows the same correlations with more demanding fixed effects. Depending on the specification, there are seller-variety and/or buyer-variety-country fixed effects.

Unlike the discount elasticity, the scale-free price $\tilde{p}$ is significantly correlated with bilateral market shares and seller concentration. Market power shifts the *intercept* of the pricing schedule but does not alter its *slope*. Large firms may negotiate lower base prices, but the marginal incentive regarding shipment size is determined by common technological factors.

Table 9: Scale-Free Price Level Variation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Dependent Variable					$\log(\tilde{p})$				
Independent Variable	Related Party	$\log(q)$ Relationship	$\log(q)$ Seller	$\log(q)$ Buyer	Output Share	Input Share	log(Rel) Length	log(Sellers)	log(Buyers)
Panel A: Baseline Level Effects									
Coefficient	0.0555	-0.172	-0.137	-0.0877	-0.154	-0.125	0.000909	0.0175	0.033
	(0.00437)	(0.00292)	(0.00238)	(0.00178)	(0.00555)	(0.00538)	(0.00153)	(0.00270)	(0.00411)
R^2	0.0863	0.217	0.177	0.133	0.0887	0.0878	0.086	0.0863	0.0864
Within R^2	0.000425	0.144	0.0995	0.0514	0.003	0.00206	0.000000706	0.000429	0.000546
Fixed Effects									
Product-Country	✓	✓	✓	✓	✓	✓	✓	✓	✓
Panel B: Complex Fixed Effects									
Coefficient	-0.011	-0.0289	-0.0501	-0.0281	-0.131	-0.0454	-0.000676	-0.0169	0.0162
	(0.00250)	(0.00152)	(0.000860)	(0.000724)	(0.00646)	(0.00516)	(0.00254)	(0.00140)	(0.00140)
R^2	0.854	0.855	0.686	0.749	0.854	0.854	0.854	0.747	0.68
Within R^2	0.000123	0.009	0.0189	0.00732	0.00342	0.000866	0.000000189	0.000655	0.0002
Fixed Effects									
Seller-Variety	✓	✓			✓	✓		✓	
Buyer-Variety-Country	✓	✓	✓	✓	✓	✓	✓		✓

Notes: This table reports regressions of the scale-free price level on various relationship and market characteristics. Panel A includes product-country fixed effects. Panel B includes more complex fixed effects, including seller-variety and buyer-variety-country fixed effects. Standard errors are in parentheses.

6.2 From Micro-Scale to Macro-Pass-Through

Our analysis reveals a stark divergence between reduced-form and structural interpretations of price data. At the micro-level, we find that "scale economies" are largely a phenomenon of shipment size rather than firm size. At the macro-level, this distinction is critical for understanding the incidence of trade shocks.

If scale economies were driven by firm-level market power or aggregate production functions, a tariff shock might be absorbed by a reduction in firm-level markups or a

movement along a firm-level cost curve. However, these scale economies can be driven by transaction-level logistics (e.g., fixed ordering costs), not firm-level production functions. As tariffs raise the holding cost of inventory, firms optimally reduce their order sizes. This reduction in transaction size forces buyers up the quantity discount schedule. Consequently, observed unit prices rise for two distinct reasons: the direct cost of the tariff, and the mechanical loss of scale economies. Standard pass-through estimates conflate these two effects, mistakenly attributing the mechanical price increase to strategic pass-through. In the next section, we show that once we correct for this mechanical "scale effect," the true pass-through of the 2018 US tariffs is significantly lower than previously thought.

7 Tariff Pass-Through

The finding that quantity discounts are cost-driven ($\gamma \approx \gamma^c$) fundamentally alters the interpretation of aggregate price movements. In standard models, incomplete pass-through implies markup compression. Our results suggest a countervailing force: if tariffs reduce trade volumes, they push firms up their average cost curves, mechanically raising unit prices.

This implies that standard estimates of unit-value pass-through do not alone measure incidence. A tariff-induced price increase could reflect two distinct phenomena: the tax itself, or the loss of scale economies due to shrinking transaction sizes. By failing to account for the latter, we risk underestimating if foreign firms are actually cutting markups (or scale-free prices) to maintain market share, hiding a substantial welfare loss borne by foreign exporters.

7.1 2018-2019 US Tariffs

The 2018-2019 trade war provides a quasi-experiment for studying tariff pass-through. During this period, the United States imposed substantial tariffs on a wide range of imported goods, with a primary focus on products originating from China. The aggregate impact of these tariffs is evident in Table 10. Total duties collected on US imports more than doubled from $29.1 billion in 2017 to $64.8 billion in 2019, raising the average applied tariff rate from 1.6% to 2.9%. The burden fell disproportionately on imports from China, where the average tariff rate nearly tripled from 3.2% to 9.5% over the same period. In contrast, the average tariff rate on imports from the rest of the world remained stable at around 1.3%.

Table 10: Summary of US Import Data

Origin	Year	Value (FOB)	Duty	Shipping Charges	Tariff Rate	Shipping Rate
		(Billion USD)			(Applied %)	(Effective %)
All	2017	1,796	29.1	52.0	1.6%	2.9%
All	2018	2,005	42.9	57.2	2.1%	2.9%
All	2019	2,274	64.8	63.2	2.9%	2.8%
All	2020	2,316	56.5	67.1	2.4%	2.9%
China	2017	363	11.6	14.3	3.2%	3.9%
China	2018	398	19.6	16.5	4.9%	4.1%
China	2019	425	40.2	17.4	9.5%	4.1%
China	2020	460	37.0	20.5	8.0%	4.5%
Other	2017	1,434	17.5	37.7	1.2%	2.6%
Other	2018	1,606	23.3	40.7	1.5%	2.5%
Other	2019	1,848	24.6	45.8	1.3%	2.5%
Other	2020	1,855	19.5	46.6	1.1%	2.5%

Notes: Data are from the US Census LFTTD and cover US imports exceeding the de minimis threshold. Value (FOB) is the free-on-board value of imports. Duty is the total customs duties paid. Shipping Charges include freight and insurance costs. The Tariff Rate is the applied rate, calculated as Duty divided by Value. The Shipping Rate is calculated as Shipping Charges divided by Value.

The standard approach to measuring tariff pass-through treats observed unit price as the relevant object for welfare analysis. However, when transaction-level quantity discounts are pervasive, this approach conflates two distinct effects: changes in the scale-free price schedule $\tilde{p}$ and compositional shifts in transaction sizes. Our framework from Section 2.2 clarifies this distinction. The total change in observed prices reflects both the shift in the underlying cost and markup structure (captured by $d\tilde{p}/d(1+\tau)$) and the adjustment in transaction quantities (captured by $\gamma \cdot dq/d(1+\tau)$). Moreover, to measure aggregate pass-through, we must aggregate transaction-level prices to construct unit values, accounting for both the intensive margin (changes in the size of existing transactions) and the extensive margin (entry and exit of trading relationships). Only by separating these components can we accurately measure the division of the tariff burden between buyers and sellers.

This distinction matters fundamentally for incidence. As we showed in our conceptual framework, the change in buyer surplus depends on shifts in the scale-free price $\tilde{p}$ weighted by quantities purchased (Equation 2), while changes in producer surplus require decomposing both revenue and cost effects (Equation 3). When transaction sizes decline in response to tariffs—as our data show they do—standard measures of pass-through will overstate the extent to which sellers can shift the tariff burden to buyers. This is because smaller transactions occur at higher unit prices due to scale economies, creating the appearance of greater pass-through even when the underlying price sched-

ule has shifted down substantially.

Our empirical strategy leverages the transaction-level variation we document to construct scale-free price measures (an implicit unit-value price index). By recovering the scale elasticity γ from within-relationship variation (Section 4.2) and using it to compute scale-free prices $\tilde{p}$, we isolate the true shift in the price schedule from compositional effects. This allows us to decompose the incidence of the 2018 tariffs into three components: the mechanical effect of the tariff itself, the adjustment in seller markups and costs (captured by changes in $\tilde{p}$), and the quantity response both at the transaction level and through entry and exit. Our findings reveal that accounting for transaction-level scale economies reduces estimated pass-through from near unity to approximately 60 percent, implying that foreign exporters absorbed a substantially larger share of the tariff burden than conventional estimates suggest.[24]

Essentially, we empirically differentiate reduced-form versus structural estimates of pass-through (MacKay, Miller, Remer, and Sheu, 2014). We do find that at both the aggregate and transaction level, we can match the reduced-form findings of the literature (Fajgelbaum et al., 2020). Considering an aggregate variety, US tariffs on imports exhibit close to full reduced-form pass-through of tariffs to unit wholesale prices. This not only is at the aggregate-level, but also at the bilateral buyer-seller-variety level (Handley, Kamal, and Monarch, 2025). However there are large quantity responses, not just from the aggregate demand curve, but also in the size of individual transactions.

Empirical Framework for Tariff Pass-Through with Transaction-Level Scale Economies. At the transaction level, we start off with the tariff-inclusive transaction price within a relationship:

$$\log p_{i\to j,t,v} = \underbrace{\log\left(1+\tau_{t,v}\right) + \log \mu_{i\to j,t,v} + \log \tilde{c}_{i\to j,t,v}}_{\log \tilde{p}_{t,v}} + \gamma \log q. \tag{16}$$

Prices are a combination of tariffs (ad-valorem), scale-free markups and marginal costs, and a final scale effect. From the analysis in Section 3, we show that this scale effect is driven by underlying costs and we will consider them invariant to tariff levels.[25]

For any given transaction t from seller i to buyer j at date d, as shown in section 4.2, there is a scale pricing elasticity γ. When considering tariffs, tariffs have log-separable

[24]Unlike our main regressions where prices were demeaned, here we keep them at their initial levels. Additionally, for the aggregate regressions, we do not drop 'singletons.'

[25]Table 15 shows that γ is invariant to tariffs.

effects on duty-inclusive prices:

$$\frac{\partial \log p_{i\to j,t,v}(q)}{\partial \log(1+\tau_{t,v})} = \underbrace{1 + \frac{\partial\left[\log \mu_{i\to j,t,v} + \log \tilde{c}_{i\to j,t,v}\right]}{\partial \log(1+\tau_{t,v})}}_{\partial \tilde{p}/\partial(1+\tau)} + \gamma \frac{\partial \log q}{\partial \log(1+\tau_{t,v})}. \tag{17}$$

The first element on the right-hand side is the effect of a tariff on a scale-independent price. This can be further decomposed into a markup effect μ, and a scale-free cost effect $\tilde{c}$. The second element is the effect of a tariff on the scale of an individual transaction.[26]

Table 11: Transaction-level Tariff Pass-Through: Within Relationships

	(1) $\log(p)$	(2) $\log(\tilde{p}^{OLS})$	(3) $\log(\tilde{p}^{IV})$	(4) $\log(pq)$	(5) $\log(q)$	(6) $\log(p)$
log(1+ Tariffs Applied)	0.956	0.780	0.800	0.357	-0.600	
	(0.0424)	(0.0282)	(0.0360)	(0.0438)	(0.0603)	
log(1+ Tariffs Statutory)						0.920
						(0.0417)
R^2	0.965	0.967	0.996	0.725	0.912	0.965
Within R^2	0.000746	0.000809	0.000358	0.0000486	0.000103	0.00045
Fixed Effects	Buyer-Seller-Variety, Variety-Year-Month					

Notes: This table reports transaction-level duty-inclusive tariff pass-through estimates for continuing buyer-seller-variety relationships in 2018. All specifications include buyer-seller-variety and variety-year-month fixed effects. Standard errors are clustered at the relationship level. Column (1) shows the pass-through of applied tariffs (duty paid) to observed transaction prices. Columns (2) and (3) report pass-through to scale-free prices ($\tilde{p}$), constructed using the OLS and IV estimates of the scale elasticity γ from Section 4.2, respectively. These columns isolate the change in the price schedule from compositional effects. Column (4) reports the effect on transaction value, while Column (5) shows the response of transaction quantity. Column (6) shows the pass-through of statutory tariffs.

7.1.1 Within-Relationship Transaction-Level Pass-Through

Looking at data on tariffs imposed by the US on its own imports, Table 11 shows the results of considering Equation (17). We consider prices, not only within a buyer-seller-variety relationship, but relative to variety-time trends (to partially account for the supply shock effect in Gertler (2023) and the diversions seen in Fajgelbaum, Goldberg, Kennedy, Khandelwal, and Taglioni (2024) and Farrokhi and Soderbery (2020).

Column (1) shows that the pass-through of tariffs to observed transaction prices in 2018 is nearly complete, with a coefficient of 0.956. This finding, which aligns with the

[26]Decomposing incomplete pass-through—distinguishing between costs, markup adjustment, and mechanical invoicing effects—is a focus of the exchange rate literature (Goldberg and Hellerstein, 2008).

reduced-form results in the literature, suggests that for every 10% increase in the tariff rate, the post-tax price paid by the buyer increases by almost 10%. However, this masks a significant quantity response. Column (5) reveals that a 10% tariff rate increase leads to a 6% decrease in the size of individual transactions.[27]

Our framework allows us to disentangle these effects. Columns (2) and (3) report the pass-through to the scale-free price, $\tilde{p}$, which isolates the shift in the underlying price schedule. Using our preferred IV estimate for the scale elasticity, we find that the pass-through to the scale-free price is only 0.80 (Column 3). This implies that the composite of seller markups and costs fell by 20% in response to the tariff. If we assume that marginal costs for a standardized order size did not decrease, this suggests that foreign exporters absorbed a significant portion of the tariff burden by reducing their markups. Had transaction sizes not fallen, the observed pass-through would have been only 80%, below the near-unity estimates that ignore quantity adjustments. This highlights a key difference from frameworks like Fajgelbaum et al. (2020), which, by assuming perfect competition and no quantity discounts ($\gamma = 0$), would interpret the full price increase as being borne by US buyers.

The decomposition in Equation (17) provides a clear interpretation of these results. The coefficient in Column (1) represents the total pass-through to the observed price, $\partial \log p / \partial \log(1+\tau)$. The coefficient in Column (3) is the pass-through to the scale-free price, $\partial \log \tilde{p} / \partial \log(1+\tau)$. The difference between these two, $0.956 - 0.800 = 0.156$, represents the quantity adjustment effect, $\gamma \cdot \partial \log q / \partial \log(1+\tau)$. Using our estimated γ of -0.29 and the quantity response from Column (5) of -0.604, this effect is $(-0.29) \times (-0.600) = 0.174$, which closely matches the observed difference. Furthermore, the pass-through to the scale-free price can be decomposed into the mechanical tariff effect (1) and the change in the composite of seller markups and costs. The estimate of 0.8 implies that this composite fell by 20 percent ($0.8 - 1 = -0.2$), indicating that foreign exporters absorbed a substantial portion of the tariff by reducing their pre-tax prices.

However, these results are based on continuing relationships and thus only capture the intensive margin of adjustment among stable relationships. Essentially, this is pass-through conditional on survival and constant order frequency and not an estimate of aggregate pass-through. To assess the full welfare incidence of the tariffs, we must aggregate these transaction-level findings to account for compositional changes.

[27] In Appendix B.4, we replicate this analysis using the tariffs imposed by the US on consumer goods in the last quarter of 2019, finding similar results. However, the estimates are noisier due to the shorter time frame before the onset of the COVID-19 pandemic.

7.1.2 Aggregate Pass-through

Moving beyond transaction-level pass-through, we now examine how quantity discounts affect aggregate unit values—the standard object used in prior tariff incidence studies. Aggregation introduces two key margins of adjustment: the intensive margin (changes in transactions and transaction sizes within and across continuing relationships) and the extensive margin (entry and exit of trading relationships). To properly measure aggregate pass-through, we must account for both how the scale-free price $\tilde{p}$ shifts in response to tariffs and how the composition of transactions changes. We present two complementary approaches: first, a scale-free aggregation that uses our estimated scale elasticity γ to construct standardized unit values from transaction-level data; and second, a first-order approximation using only aggregate data that yields quantitatively similar results and can be implemented without access to microdata.

Next, we aggregate these scale-free prices across all transactions for a given product-origin-date combination, weighting by transaction quantities to obtain an aggregate scale-free unit value $\tilde{p}_{o,d,v}$, from origin country o to the destination d (the US in our case) for variety v. This approach allows us to isolate the shift in the underlying price schedule from compositional changes in transaction sizes and counts.

To construct scale-free unit values, we follow two steps. We first adjust transaction-level prices to remove the effect of quantity discounts using our estimated scale elasticity γ_v. This yields a scale-free price $\tilde{p}_{i\to j,t,v}$ for each transaction that reflects what the price would be if the transaction size were one unit:

$$\tilde{p}_{i\to j,t,v} = \frac{p_{i\to j,t,v}}{q_{i\to j,t,v}^{\gamma_v}}, \tag{18}$$

where γ is recovered through the mechanism in section 4.2.[28]

We create a duty-inclusive unit value in the spirit of aggregate regressions in Fajgelbaum et al. (2020); Amiti et al. (2019), who compute:

$$p_{o,d,v} = \sum_{i\in o,t\in d} p_{i\to j,t,v} q_{i\to j,t,v} / \sum_{i\in o,t\in d} q_{i\to j,t,v}. \tag{19}$$

Our scale-free prices aggregates across all buyers and sellers, where shipment of variety

[28] Alternatively, we can use $\gamma_{i\to j,v}$ for the same effect.

v originates in country o at date d, aggregated at the monthly level:[29]

$$\tilde{p}_{o,d,v} = \sum_{i \in o, t \in d} \tilde{p}_{i \to j,t,v} q_{i \to j,t,v} / \sum_{i \in o, t \in d} q_{i \to j,t,v}. \tag{20}$$

With this approach, we can decompose the duty-inclusive aggregate price for a product as:

$$\log p_{o,d,v} = \log \tilde{p}_{o,d,v} + \log \tilde{q}_{o,d,v}, \tag{21}$$

where the residual $\tilde{q}_{o,d,v}$ captures the aggregate quantity composition effect. This term reflects how the distribution of transaction sizes affects the observed unit value, even holding the underlying price schedule $\tilde{p}$ fixed. The scale-free price $\tilde{p}_{o,d,v}$ can be interpreted as the tariff-inclusive composite of scale-free markups and marginal costs: $\tilde{p}_{o,d,v} = (1 + \tau_{o,d,v})(\mu\tilde{c}_{o,d,v})$, where $\mu\tilde{c}_{o,d,v}$ represents the quantity-weighted average of seller markups and costs, net of tariff charges.

Table 12 regresses each of these components, both the left- and right-hand sides, on changes in applied tariffs. We consider different levels of aggregation, the baseline in Panel (A) considers just monthly 2018 data. We later consider robustness in Panels (B)-(E), aggregating up to the yearly level.

Column (1) replicates the baseline findings, pass-through, not just at the transaction level, but the product-origin level is effectively unity. Columns (2) and (3) adjust for quantity discounts using our OLS and IV estimates for the scale elasticity γ_v within variety.

We can further use the difference between the aggregate un-adjusted regression in column (1) with the scale-free results in columns (2)-(5) to consider the relative effects on sellers. In our baseline comparing to our IV results, a 10% increase in tariffs decreases the prices received by sellers by 4.0%. Without the scale adjustment, the prices would have increased by 0.5%

The difference between the unweighted transaction-level estimate (0.80) and the volume-weighted aggregate estimate (0.60) indicates that compositional effects are significant. Larger transactions, which benefit more from quantity discounts, tend to shrink more in response to tariffs. This compositional shift amplifies the mechanical price increase due to reduced transaction sizes, leading to a lower aggregate pass-through rate.

The microdata founded scale elasticities play a crucial role. But how important is microdata for this exercise? As an alternative, we bound our results using a much more straightforward decomposition, conditional on recovering scale elasticities from the microdata.

[29] Using expenditure weights does not yield a significant quantitative difference.

As an alternative to Equation (21), we can interpret results using a much more straightforward decomposition,

$$\log p_{o,d,v} = \log pq_{o,d,v} - \log q/T_{o,d,v} - \log T_{o,d,v}, \tag{22}$$

where T denotes the number of transactions. Here, we decompose the aggregate unit value into three components: the total payments received by sellers for variety v from origin o at date d (i.e., total revenue) denoted by $pq_{o,d,v}$, the average order size captured by $q/T_{o,d,v}$, and the number of transactions $T_{o,d,v}$.

Table 12 in columns (5-8) regresses tariff rates on each of these components. Column (5) indicates a 10% increase in tariffs leads to an 8.7% decrease in payments. Column (7) shows that a 10% increase in tariffs leads to a 15.5% decrease in average order size. Column (8) shows that a 10% increase in tariffs leads to a 3.6% decrease in the number of transactions.

In particular the second term, $\log q/T_{o,d,v}$, is informative of the level of pass-through netting out quantity changes, in combination with our γ_v from Section 4.2. As an approximation, if we assume the change in average order size is reflective of the weighted average change in order size at the transaction level, we can use our earlier estimate of γ to back out the change in scale-free prices:

$$\frac{d\log \tilde{p}_{o,d,v}}{d\log(1+\tau_{o,d,v})} \approx \frac{d\log p_{o,d,v}}{d\log(1+\tau_{o,d,v})} - \gamma\frac{d\log\left(q/T\right)_{o,d,v}}{d\log(1+\tau_{o,d,v})} = 1.05 - (-1.55) * (-0.29) = 0.60.$$

Essentially, recovering columns (2)-(4) to a first order approximation. For a consistent order size, tariffs are only passed through at the aggregate level at 60%.[30] As with before, we can decompose this into a tariff direct effect (mechanically 1) and a markup and cost effect. If we assume that costs did not fall during the tariff war, then aggregate markups fell 40%.

The discrepancy between the relationship-level pass-through (Table 11) and the aggregate results stems from the extensive margin. As shown in Column (7), the number of transactions declines significantly in response to tariffs. This reduction in order frequency, combined with the reduction in order size, compounds the welfare loss. While continuing relationships see a pass-through of roughly 0.80 to scale-free prices, the exit of relationships and the consolidation of orders implies that the aggregate average unit value rises by more than the intensive-margin estimate would suggest. This aligns with work em-

[30]Due to Jensen's inequality, the logarithm of a sum is not the sum of a logarithm (as log is a concave function), thus requiring the use of detailed transaction-level data.

phasizing the extensive margin of global sourcing, where fixed costs of importing lead to sensitivity in the measure of active suppliers (Antras, Fort, and Tintelnot, 2017).

Correcting for Quality For robustness, we also consider a Feenstra-style CES price index within each origin-destination-variety combination to account for the extensive margin of entry and exit. We use data on the elasticity of substitution from a structural model of supply and demand.

In Appendix A.9, we construct an exact CES price index to account for substitution bias and variety turnover (entry and exit). Using the Sato-Vartia index for continuing goods and the Feenstra adjustment for the extensive margin, we find that standard unit-value indices may overstate inflation by failing to account for substitution towards cheaper varieties. However, for the medium-run pass-through estimates in Table 12 column (4), the divergence between the CES index and our simple average is minimal, suggesting that substitution bias is not the primary driver of our results at this horizon.[31]

Event Horizons In Table 12 Panels (B)-(E), we consider different time periods and aggregation horizons. We first include all data from 2017-2019 at the monthly level (B), before aggregating to the quarterly (C) and year levels (D). Lastly we take the yearly data from 2017 and do a two-year difference to 2019 (E).

Echoing the robustness of the unadjusted pass-through in Fajgelbaum et al. (2020), our baseline results in column (2)-(4) adjusting our unit price indices for quantity discounts are largely consistent across panels, with measured pass-through between 0.4 and 0.7, with OLS scale elasticities at the higher end and the CES-adjustment at the lower end.

Decomposing Aggregate Pass-through In aggregate, pass-through is the combination of multiple forces, the pass-through of continuing firms, and the net effect of entrants and exiting relationships.

Here we focus on the medium-run difference. With most of the 2018 tariffs implemented in September 2018 (and slowly adjusted until January 2019), we consider the medium run difference between 2017 and 2019.[32]

In Table 13, we run effectively the same decomposition as in the previous subsection, but in first differences, as there are only two periods. In Panel A, we consider continuing

[31] In a companion project (Ganapati and Hottman, 2025) we develop a full demand system with a theoretically consistent price index that is entirely consistent with observed price schedules and pass-through rates. Our regressions are effectively difference in difference regressions, so our corrected unit values are relative to non-tariffed unit values.

[32] This ignores the set of consumer products tariffed in late 2019 on the eve of the COVID-19 pandemic.

Table 12: Decomposing Aggregate Tariff Pass-Through by Methodology and Time Period

	(1) $\log(P)$	(2) $\log(\tilde{p}_{\gamma\text{OLS}})$	(3) $\log(\tilde{p}_{\gamma\text{IV}})$	(4) $\log(\tilde{p}_{\text{CES}})$	(5) log(Value)	(6) $\log(Q)$	(7) $\log(Q/T)$	(8) $\log(T)$
Panel A: Monthly 2018 Data								
$\log(1+\text{Tariffs})$	1.047	0.804	0.604	0.573	-0.867	-1.914	-1.551	-0.362
	(0.0546)	(0.0414)	(0.0447)	(0.0418)	(0.0709)	(0.0897)	(0.0755)	(0.0394)
r^2	0.919	0.942	0.951	0.951	0.876	0.888	0.913	0.918
Within R^2	0.000478	0.000312	0.000223	0.000223	0.000155	0.000301	0.000757	0.000155
Panel B: Monthly 2017-2019 Data								
$\log(1+\text{Tariffs})$	1.038	0.547	0.537	0.504	-1.479	-2.517	-1.697	-0.820
	(0.0302)	(0.0226)	(0.0238)	(0.0466)	(0.0415)	(0.0526)	(0.0424)	(0.0258)
r^2	0.897	0.927	0.938	0.686	0.861	0.891	0.88	0.897
Within R^2	0.00065	0.000357	0.000278	0.000276	0.00125	0.00182	0.00111	0.00112
Panel C: Quarterly 2017-2019 Data								
$\log(1+\text{Tariffs})$	1.069	0.584	0.568	0.436	-1.669	-2.739	-1.751	-0.988
	(0.0362)	(0.0276)	(0.0299)	(0.0260)	(0.0507)	(0.0639)	(0.0503)	(0.0318)
r^2	0.909	0.934	0.945	0.65	0.892	0.909	0.896	0.918
Within R^2	0.000672	0.000386	0.000295	0.000456	0.00151	0.00205	0.00118	0.0014
Panel D: Yearly 2017-2019 Data								
$\log(1+\text{Tariffs})$	1.178	0.68	0.645	0.481	-2.032	-3.21	-1.886	-1.324
	(0.0578)	(0.0469)	(0.0522)	(0.0240)	(0.0830)	(0.104)	(0.0809)	(0.0528)
r^2	0.937	0.954	0.962	0.578	0.934	0.94	0.928	0.949
Within R^2	0.000808	0.000505	0.000367	0.000801	0.0021	0.00267	0.00138	0.00212
Panel E: Two Year Annual Change 2017 to 2019 Data								
$\log(1+\text{Tariffs})$	1.093	0.645	0.602	0.533	-2.189	-3.283	-1.737	-1.546
	(0.0654)	(0.0516)	(0.0591)	(0.0303)	(0.0938)	(0.118)	(0.0891)	(0.0630)
r^2	0.949	0.963	0.969	0.544	0.938	0.947	0.939	0.95
Within R^2	0.00105	0.000684	0.000471	0.00122	0.00321	0.00388	0.00172	0.00359
Fixed Effects	Variety-Country Origin, Time-Month							

Notes: This table reports aggregate tariff pass-through estimates at the origin-destination-variety level using different methodologies and time periods. Each panel represents a different time aggregation: Panel A uses monthly data from 2018, Panel B uses monthly data from 2017-2019, Panel C uses quarterly data from 2017-2019, Panel D uses yearly data from 2017-2019, and Panel E uses two-year changes from 2017 to 2019. Column (1) shows the pass-through of applied tariffs to observed unit values. Columns (2) and (3) report pass-through to scale-free unit values, constructed using the OLS and IV estimates of the scale elasticity γ_v, respectively. Column (4) presents pass-through estimates using a CES price index to account for quality adjustments. Columns (5) to (8) decompose the aggregate unit value into total value, total quantity, average order size, and number of transactions. All specifications include variety-country origin and time-month fixed effects. Standard errors are clustered at the variety-country origin level.

trading partners that account for approximately 2/3 of trading volume. In this case, aggregate pass-through is only 0.87 (Col 1), with a large order size and order number effect (Col 5-6). However, our scale-corrected pass-through estimates are .60 and .57 (Col 2-3).

Panel B considers the net effect of entrant and exiting trading pairs. Aggregate pass-through, under the strict assumption of comparability, implies that pass-through is more

than full (1.2) for this set of firms (Col 1) with even larger large order size and order number effects (Col 5-6). But these differences with continuing relationship are minimized when considering the scale-corrected pass-through estimates of .64 and .60 (Col 2-3).[33]

There are massive compositional changes induced by tariffs, but estimates correcting for quantities aligned our estimates of pass-through between various samples.

Unification With Exchange Rate Pass-Through The nearly complete pass-through of tariffs into aggregate unit import prices (Amiti et al., 2019; Fajgelbaum et al., 2020; Cavallo et al., 2021) stands in stark contrast to the Exchange Rate Pass-Through (ERPT) literature, which typically finds much lower pass-through rates into import prices, often in the range of 0.4 to 0.6, even in the medium run (Gopinath et al., 2010). This discrepancy presents a puzzle: why would foreign exporters absorb exchange rate shocks but fully pass on tariff shocks?

A key distinction between these literatures lies in the data used. Tariff studies often rely on customs unit values, which are calculated as total value divided by total quantity. As we have shown, unit values mechanically incorporate changes in order size and the resulting quantity discounts. In contrast, the ERPT literature, such as Gopinath et al. (2010), typically utilizes BLS import price indices. These indices are constructed from survey data that track the price of specific items over time, filtering out the quantity discount channel we highlight.[34]

When we correct for scale effects and isolate the "scale-corrected" pass-through—analogous to controlling for item specifics and quantity—we estimate a pass-through rate of approximately 0.60 - remarkably close to the standard ERPT estimates. The apparent difference between tariff and exchange rate pass-through may be driven by the measurement of prices: unit values capture the endogenous response of quantities and discounts, whereas price indices isolate the pure price change.

This lines up with literature that emphasizes strategic complementarities and variable markups as the primary drivers of incomplete pass-through (Atkeson and Burstein, 2008; Amiti et al., 2014). Both tariff and exchange rate shocks affect the underlying cost structure of exporters, leading to adjustments in markups. When controlling for quantity effects, both types of shocks exhibit similar pass-through patterns, suggesting that the underlying market dynamics are consistent across these different types of trade shocks.

[33]These estimates use the same scale elasticities as in Panel A, which are derived from our within-relationship IV estimates for γ_v.

[34]The BLS for domestic price indices directly account for quantities through survey data. For their international trade index, they use administrative data for 2/5 of the sample (U.S. Bureau of Labor Statistics, 2025).

Table 13: Decomposing Medium Run Tariff Pass-Through

Panel A: Continuing Relationships from 2017 to 2019

	(1) $\log(P)$	(2) $\log(\tilde{p}_{\gamma\text{OLS}})$	(3) $\log(\tilde{p}_{\gamma\text{IV}})$	(4) $\Delta\log(\text{Value})$	(5) $\Delta\log(Q)$	(6) $\Delta\log(Q/T)$	(7) $\Delta\log(T)$
$\Delta\log(\text{Tariffs})$	0.871	0.603	0.572	-0.861	-1.731	-0.73	-1.002
	(0.0607)	(0.0450)	(0.0517)	(0.0863)	(0.104)	(0.0763)	(0.0672)
R^2	0.129	0.123	0.126	0.127	0.133	0.125	0.136
Within R^2	0.00189	0.00165	0.00113	0.000916	0.00254	0.000843	0.00204

Panel B: Other Firms

	(1) $\log(P)$	(2) $\log(\tilde{p}_{\gamma\text{OLS}})$	(3) $\log(\tilde{p}_{\gamma\text{IV}})$	(4) $\Delta\log(\text{Value})$	(5) $\Delta\log(Q)$	(6) $\Delta\log(Q/T)$	(7) $\Delta\log(T)$
$\Delta\log(\text{Tariffs})$	1.222	0.637	0.6	-2.857	-4.079	-2.278	-1.801
	(0.0964)	(0.0697)	(0.0779)	(0.114)	(0.155)	(0.121)	(0.0757)
R^2	0.0887	0.0896	0.0846	0.119	0.112	0.0949	0.134
Within R^2	0.00108	0.000559	0.000397	0.00421	0.00464	0.00238	0.00378
Fixed Effects				Variety-Year			

Notes: This table decomposes medium-run tariff pass-through by relationship status. Panel A reports estimates for continuing relationships. Panel B reports estimates for the net effect of entering and exiting relationships. All specifications are in first differences between June 2018 and June 2019 and include variety-year fixed effects. Standard errors are clustered at the variety-origin level.

7.2 Revisiting Incidence

Overall, using a conventional unit price index—as opposed to one adjusted for transaction quantities—would suggest near-complete pass-through. However, our scale-free results indicate a net pass-through of approximately 50-60%, with the difference driven by substantial decreases in average order size and the number of orders.

Having established that pass-through is incomplete and quantity responses are large, we now combine these estimates to assess the overall welfare burden. We return to the framework from Section 2.2. While a full welfare analysis typically requires a structural supply and demand system, Weyl and Fabinger (2013) show that marginal effects can be recovered using sufficient statistics. We extend this logic to our setting, incorporating the imperfect competition wedges we have empirically identified.

We model seller prices p_s, revenues R, and costs C as:

$$p_s = \frac{\tilde{p}}{1+\tau} q^{\gamma_c}, \quad R = \frac{\tilde{p}}{1+\tau} q^{\gamma_c+1}, \quad C = \tilde{c} q^{\gamma_c+1}$$

Locally, the changes in revenue and cost with respect to tariffs are:

$$\frac{dR\left(q\right)}{d\left(1+\tau\right)}=\left[\frac{d\log\tilde{p}}{d\log\left(1+\tau\right)}+\left(\gamma_{c}+1\right)\frac{d\log q}{d\log\left(1+\tau\right)}-1\right]\frac{R}{1+\tau}$$

$$\frac{dC\left(q\right)}{d\left(1+\tau\right)}=\left[\frac{d\log\tilde{c}}{d\log\left(1+\tau\right)}+\left(\gamma_{c}+1\right)\frac{d\log q}{d\log\left(1+\tau\right)}\right]\frac{R}{\tilde{\mu}\left(1+\tau\right)}$$

Note that terms involving γ^{μ} drop out due to our finding that quantity discounts are cost-driven ($\gamma^{\mu}\approx 0$). To complete this calculation, we need two additional pieces of information: the elasticity of scale-free costs $\tilde{c}$ to tariffs, and the baseline level of markups $\tilde{\mu}$. In addition, we need to test the implicit assumption that γ is invariant to tariffs.

Elasticity of Scale-Free Cost to Tariffs Understanding how scale-free costs change due to the trade war is not trivial. In principle, this needs to be known for every existing relationship. As we are primarily interested in incidence, we proceed under the assumption that the level of marginal costs is invariant to the destination, i.e., $\tilde{c}_{odv}=\tilde{c}_{ov}$, and that markups to foreign markets are unaffected.

We assume that products sold to the United States have the same cost structure as those sold to the rest of the world and that markups to 3rd parties are also unchanged. Essentially, the idea is that goods sold to Europe and America use the same mix of inputs. We test this by examining unit prices of exports to third-party countries:

$$\frac{d\log\tilde{p}_{o,-US,v}}{d\log(1+\tau_{o,US,v})}=\frac{d\log\tilde{c}_{o,US,v}}{d(1+\tau_{o,US,v})},$$

where we look at shipments from origin o to non-US destinations $-US$ for variety v. As we do not have transaction-level data, we leverage results from the prior section and consider using aggregate shipment quantities to control for scale-effects.

Table 14 highlights the recovery of the pass-through of US tariffs on goods exported via third party trade, controlling for the transaction-level aggregation in Section 6. We consider data from 2017 (pre-tariff) and 2019 (post-tariff). Columns (1) and (2) conduct a first difference regression with varying levels of fixed effects and controlling for aggregate scale economies. We find that US tariffs did not change $\tilde{p}$ for third countries. Column (3) conducts a two-stage regression, first netting out scale economies and then analyzing the residual price $\tilde{p}$. We find no significant effect. This empirical result is extremely tight around zero:

$$\frac{d\log\tilde{c}_{o,US,v}}{d\log(1+\tau_{o,US,v})}\approx 0.$$

Table 14: Decomposing US Tariff Pass-Through on Foreign Prices

	(1) $\Delta \log \tilde{p}_{i,j\neq US,v}$	(2) $\Delta \log \tilde{p}_{i,j\neq US,v}$	(3) $\Delta \log \tilde{p}_{i,j\neq US,v}$
$\Delta \log(1+\tau_{i,US,v})$	-0.0171	0.00758	-0.0193
	(0.0200)	(0.0108)	(0.0227)
N	4,847,859	1,552,523	355,727
Fixed Effects			
Destination-Variety	✓		
Origin-Destination	✓		
Sample	All	$\Delta\tau \neq 0$	Chinese Origin

Notes: This table estimates the pass-through of US tariffs on scale-free prices to third-party countries. We report first-difference OLS estimates with varying fixed effects and controls for scale economies. We perform a two-stage least squares estimates, first netting out scale economies and then regressing the residual scale-free price on US tariffs. Column (1) includes destination-variety and origin-destination fixed effects. Column (2) only considers products with changes in the US-origin Import tariff rate. Column (3) only considers products originating in China. Standard errors are clustered at the product-origin level.

Markups To complete our incidence calculation, we require an estimate for the level of markups, $\tilde{\mu}$. We compute our own markup estimates using the methodology of Feenstra (1994) as detailed in Appendix B.3. This approach assumes CES demand and monopolistic competition to recover markups from trade data.[35] Our secondary estimates are drawn from the literature, specifically studies using production data from China which find modal markups around 1.3 for exporters (Yue and Lin, 2023).

Scale Changes Implicit in our analysis, we have assumed that the scale elasticity $\gamma_{o,v}$ is constant across all products and origins and invariant to tariffs. We test this assumption by estimating changes in $\gamma_{o,v,year}$ directly from the data before and after the 2018 tariffs. We retain the OLS framework from Table 4 and allow γ to vary by origin, variety, and year.[36] Table 15 shows that the changes in the scale elasticity is invariant to tariff changes from 2017 to 2019. Column (1) is the baseline regression in first differences. Columns (2-4) conduct robustness and controls for other factors such as country and HTS6 codes

[35]It is important to note that this method is only locally valid in our framework and is globally inconsistent with our main findings for two reasons. First, a standard CES monopolistic competition model cannot rationalize our observed aggregate pass-through of less than one in the presence of a downward-sloping supply curve; such a model would predict pass-through greater than one. Second, our main analysis is more general and does not impose the theoretical consistency condition that the demand elasticity must be greater than the supply-side scale elasticity. While this condition is necessary for equilibrium stability in a fully specified structural model, our reduced-form approach allows us to estimate these elasticities separately.

[36]Using the IV is problematic here, with no first stage power.

Table 15: Transaction Scale Changes and Tariff Pass-Through

	(1)	(2)	(3)	(4)
	2017-2019 Change Pricing Scale: $\Delta\gamma^{OLS}_{o,d,v}$			
Δ(1+Tariff)	-0.0264	-0.00581	0.181	0.0127
	(0.0567)	(0.0102)	(0.123)	(0.131)
Fixed Effects			Country	HTS6
Weighted	Observations	Quantile Regression	Observations	Observations

Notes: This table reports regressions of the change in the estimated quantity discount parameter ($\Delta\gamma$) between 2017 and 2019 on the change in tariffs over the same period. The dependent variable is the change in the OLS-estimated scale parameter from Table 4 recovered at the level of origin-destination-variety-year. Column (1) is a simple OLS regression weighted by the number of transactions. Column (2) reports a quantile regression at the median. Columns (3) and (4) include country and HTS6 fixed effects, respectively. Standard errors are in parentheses.

or considering quantile regression.

Welfare Implications Table 16 presents the welfare decomposition. We approximate the total change in surplus by applying these marginal incidence rates to the observed change in tariff revenue and trade volumes. Row 1 shows our baseline estimate: a $22.9 billion reduction in foreign producer surplus versus a $16.6 billion loss for domestic buyers. This implies an incidence ratio (I) of 0.72, a reversal from non-scale adjusted estimates where buyers bear the burden (incidence $I \geq 1$).

This finding helps discern the underlying market structure. At the bottom of Table 16, we report theoretical benchmarks derived from Section 2. Under standard models like Perfect Competition or Symmetric Oligopoly, incidence is dominated by domestic buyers (Incidence ratios of ∞ or close to 1, with foreign incidence near zero). Our baseline result of $I = 0.72$ is incompatible with these standard frameworks. Instead, it lies closer to the "Perfectly Discriminatory Monopoly" benchmark ($I = 0$), suggesting that the pervasive quantity discounts we observe act as a mechanism for exporters to absorb tariffs by effectively pricing to market.

Comparing the "Baseline" (Row 1) to "Rest of World" (Row 3), we see the mechanism in action. For non-China trade—where tariffs were lower and supply chains less disrupted—the incidence ratio (0.85) is higher and closer to the oligopoly benchmark. In contrast, the US-China trade war (Row 2) drove incidence down to 0.71. This suggests that when faced with large, persistent tariffs, Chinese exporters utilized the quantity-discount margin aggressively to maintain market access, slashing effective markups more

than their counterparts elsewhere.

In aggregate, the tariffs functioned as a tool of international redistribution, but with significant efficiency costs. Our calculation in Row 1 suggests a net gain for the US (Buyer Loss + Govt Revenue $\approx$ \$13.4 billion gain). This result relies on our finding that the relevant price for Buyer Surplus is the scale-free price $\tilde{p}$, not the observed unit value. While observed unit values rose one-for-one with tariffs, a significant portion of this increase was due to a mechanical slide up the cost curve as buyers reduced order sizes. By the Envelope Theorem, the welfare loss from this quantity adjustment is second-order; the primary distortion is the shift in the price schedule itself, which rose by only 60%.

However, this domestic neutrality came at the cost of a massive \$22.9 billion profit shock to foreign producers. This figure captures both the reduction in markups and the true "efficiency loss" of the trade war—the real resources consumed by processing smaller, more frequent batches. Because transaction-level costs rose as scale economies were lost, foreign sellers faced a double squeeze: lower effective prices and higher per-unit costs., consistent with the production declines documented in Chor and Li (2024).[37]

It is important to qualify that the Envelope Theorem provides a first-order approximation of welfare changes. While standard (Weyl and Fabinger, 2013; Ganapati et al., 2020), our estimated quantity responses suggest that second-order effects ("Harberger triangles") could be non-negligible. Specifically, if the disutility of frequent ordering rises non-linearly, the approximation might understate the administrative burden on domestic buyers. However, even if we were to treat the loss of scale entirely as a direct cost to buyers, the qualitative failure of the "complete pass-through" remains, the reduction in $\tilde{p}$ provides a substantial transfer.

Extensions and Robustness These tariff burden results are robust across a range of alternative specifications. Rows 4-12 of Table 16 show that using different estimates for scale economies yield broadly similar results, with the seller's share of the burden remaining substantial and incidence ratios between 0.6 and 0.85. Row 4 uses OLS-derived γ_v and Row 5 uses the IV-derived common γ from Section 4. Row 6 uses the same mean $\mu = 1.3$ across industries, Row 7 uses estimates of μ from Yue and Lin (2023)'s cost based approach from Chinese microdata, and Row 8 uses $\mu = 1.5$. Row 9 considers the CES price index, instead of the baseline quantity weighting. Row 10 uses the longer-run changes from 2017-2019, instead of the baseline monthly variation (referring to Panel E in Table 12).

The main outlier is the "Transaction-Level Pass-Through" estimate (Row 11), which

[37] A caveat is that we do not attribute profits of multinationals to ultimate owners.

suggests buyers bore a much larger burden. This is because it is calculated using only intensive-margin adjustments within continuing relationships, thereby missing the crucial effects of firm entry and exit and other compositional shifts that are captured in our aggregate analysis.

One potential driver of these findings may be simply that $\gamma_\mu = 0$. To test the quantitative significance of this, we instead assume that scale economies are entirely driven by markups ($\gamma_c = 0$). In Row 12, we re-calculate incidence under this assumption and find similar results, with sellers still bearing a large share of the burden (Incidence = 0.64). Producers would have lost even more surplus and the burden on them would have increased.

Our baseline welfare calculations primarily capture small changes. However, as shown in Table 12, the data reveal a significant extensive margin response, with a large change in transaction frequency. In Appendix A.6, we generalize our incidence framework to explicitly account for these extensive margin effects. As of now, we are unable to directly compare our local result to the general-equilibrium result in Fajgelbaum et al. (2020).

However, by locally approximating the surplus lost from exiting transactions in Row 13, we show that the total welfare impact is larger than the intensive-margin estimates suggest and an order of magnitude larger than Fajgelbaum et al. (2020). While the qualitative distribution of the burden remains similar, there is now much larger aggregate efficiency loss, as the extensive margin changes represent a true loss of mutually beneficial trade relationships, domestic welfare now declines by $30.3 billion, outweighing government revenue. We leave a more precise comparison to Ganapati and Hottman (2025), where we fully estimate a consistent demand system.

7.3 Rationalizing the Pass-Through Mechanism

How do we reconcile our findings of incomplete base-price pass-through ($\tilde{\rho} \approx 0.6$) with the existence of significant scale economies? In a standard competitive model with constant markups, increasing returns to scale (a downward-sloping supply curve) implies that tariff pass-through should exceed 100 percent. Intuitively, because tariffs reduce the quantity demanded, and lower quantities drive up marginal costs (by moving firms up the cost curve), the price increase should be amplified beyond the tariff itself.

To rationalize these findings, we interpret them through a structural framework that accounts for simultaneous quantity discounts, imperfect competition, and variable markups. We again adapt the sufficient statistic approach of Weyl and Fabinger (2013) to our setting

Table 16: Welfare Effects of Tariffs

	Specification	Change in Surplus (Billion 2017 USD)				Incidence (Ratio)	Share of Surplus Change	
		Seller	Buyer	Gov.	Total	Buyer / Seller	Seller vs Buyer	Domestic vs Foreign
1	Baseline	-22.9	-16.6	30.0	-9.5	0.72	58%	-141%
2	Baseline - China Only	-20.9	-14.9	26.8	-9.0	0.71	58%	-132%
3	Baseline - Rest of World	-2.0	-1.7	3.2	-0.5	0.85	54%	-300%
4	OLS Scale Discounts	-26.9	-16.6	30.0	-13.5	0.62	62%	-99%
5	Common Scale Discounts	-27.0	-16.6	30.0	-13.6	0.61	62%	-99%
6	Common Markups	-19.7	-16.6	30.0	-6.3	0.84	54%	-213%
7	Exogenous Markups	-22.1	-16.6	30.0	-8.7	0.75	57%	-154%
8	Common High Markups	-24.7	-16.6	30.0	-11.3	0.67	60%	-119%
9	CES Price Index	-21.4	-18.2	30.1	-9.5	0.85	54%	-125%
10	2017-2019 Only	-22.1	-16.6	30.0	-8.7	0.75	57%	-154%
11	Transaction Level Pass-Through	-9.9	-23.7	30.6	-3.0	2.39	30%	-228%
12	Markup Driven Scale	-25.9	-16.6	30.0	-12.5	0.64	61%	-107%
13	Extensive Margin	-35.9	-30.3	30.0	-36.2	0.84	54%	1%
14	Perfect Competition					-∞	-1%	-
15	One-Price Monopoly					1.01	50%	1%
16	Perfectly Discriminatory Monopoly					0.00	100%	-∞

Notes: This table presents the welfare effects of US tariffs under various specifications. The first section (Rows 1-13) shows our baseline and robustness checks, reporting changes in foreign seller surplus, domestic buyer surplus, government revenue, total welfare change, incidence ratios, and shares of surplus change. Total represents the sum of Seller, Buyer, and Government revenue changes. The second section (Rows 14-17) provides theoretical benchmarks for incidence under different market structures. See text for details on each specification.

with ad-valorem tariffs, following Adachi and Fabinger (2022).[38] The unit-price pass-through elasticity ρ is determined by the interaction of demand, scale, and conduct:

$$\rho \equiv \frac{d\log P}{d\log(1+\tau)} = \left[\underbrace{\frac{\epsilon_D - \theta}{\epsilon_D}}_{\text{Stabilizer}}\right]\left[1 + \underbrace{\frac{\theta}{\epsilon_{ms}}}_{\text{Curvature}} + \underbrace{\frac{\theta}{\epsilon_\theta}}_{\text{Conduct}} + \underbrace{\frac{\epsilon_D - \theta}{\epsilon_S}}_{\text{Scale Effect}}\right]^{-1} \tag{23}$$

where ϵ_D is the price elasticity of demand, ϵ_S is the supply elasticity (the inverse of the elasticity of marginal cost with respect to quantity), ϵ_{ms} is the inverse elasticity of marginal surplus which captures demand curvature, and ϵ_θ is the inverse elasticity of the conduct parameter with respect to quantity. The conduct parameter θ ranges from 0 (perfect competition) to 1 (monopoly). See Appendix A.4 for a full derivation.

This decomposition illustrates that observed pass-through is the outcome of offset-

[38] We assume that γ is invariant to tariffs, as shown in the empirical exercise. The appendix generalizes results.

ting rising costs and falling markups. The *Scale Effect* term ($\frac{\epsilon_D - \theta}{\epsilon_S}$) is negative because the supply curve is downward-sloping ($\epsilon_S < 0$) and amplifies the pass-through rate (ρ). Economically, this reflects that tariffs reduce trade volumes, which forces transactions up the average cost curve, adding a mechanical cost increase on top of the tariff itself. Therefore, finding a pass-through rate that is merely unity—despite these rising costs—requires a powerful countervailing force. The remaining terms in the denominator (driven by demand curvature and conduct) must be sufficiently large and positive to stabilize prices. In short, to keep prices from spiraling upward, exporters compress their markups aggressively enough to offset the loss of scale economies.

The Scale Effect vs. The Stabilizer The decomposition highlights two opposing forces. The first is the mechanical *Scale Effect* captured by the term $\frac{\epsilon_D - \theta}{\epsilon_S}$. Because quantity discounts imply a downward-sloping supply curve ($\epsilon_S < 0$), falling volumes drive unit costs up, exerting upward pressure on pass-through ($\rho > 1$). This explains why observed unit values rose fully: tariff-induced volume collapses forced transactions up the average cost curve. Counteracting this is the leading term $\frac{\epsilon_D - \theta}{\epsilon_D}$, which acts as a *Stabilizer*. As Delipalla and Keen (1992) show, ad-valorem tariffs act as "automatic stabilizers" compared to specific taxes: because the tax payment scales with the price, firms face a stronger disincentive to raise markups. This "multiplier effect" persists even with non-linear pricing, forcing markups down to dampen the final price increase.

In imperfectly competitive markets, firms may compress markups more aggressively in response to tariffs. The finding of complete pass-through ($\rho \approx 1$) in recent trade war studies (Fajgelbaum et al., 2020; Cavallo et al., 2021) presents a paradox, as it implies either perfect competition or constant markups. Our estimate of $\tilde{\rho} \approx 0.6$ resolves this paradox, confirming that when the mechanical effects of quantity discounts are removed, the behavioral response of exporters aligns with the theoretical prediction that ad-valorem taxes are under-shifted.

Demand Curvature and Real Rigidities For the scale-free pass-through to be incomplete ($\tilde{\rho} \approx 0.6$), the "Curvature" term θ/ϵ_{ms} in the denominator must be large and positive. This parameter $1/\epsilon_{ms}$—the elasticity of marginal surplus—corresponds directly to the "super-elasticity" of demand identified by Klenow and Willis (2016) and Kimball (1995) as the source of real rigidities in macroeconomics. In trade theory, this identifies import demand as "sub-convex" or log-concave (Mrázová and Neary, 2017). Economically, this confirms "Marshall's Second Law of Demand": as tariffs push prices up, the demand elasticity ϵ_D rises, forcing firms to aggressively compress markups to defend

market share. Our results suggest that standard CES models ($1/\epsilon_{ms} = 0$) fundamentally understate these strategic complementarities.

Revisiting the "Elusive" Gains from Trade The magnitude of the pass-through denominator speaks directly to the welfare debate. Arkolakis, Costinot, Donaldson, and Rodríguez-Clare (2019) (ACDR) argue that pro-competitive gains from trade are often "elusive" because foreign markup adjustments can offset domestic benefits. In their framework, pass-through near unity (as suggested by uncorrected unit values) implies markup neutrality. However, our finding that scale-free pass-through is significantly less than unity ($\tilde{\rho} < 1$) overturns this neutrality.

We can quantify the precise magnitude of this effect using our reduced-form estimates. Exploiting the geometric relationship between unit-price pass-through ($\rho \approx 1$) and scale-free pass-through ($\tilde{\rho} \approx 0.6$), we first recover an implied demand elasticity of $\epsilon_D \approx 1.43$, consistent with differentiated goods.[39] Crucially, this allows us to isolate the elasticity of the markup ϕ_μ to a tariff shock:

$$\phi_\mu \approx \tilde{\rho} - 1 = -0.4.$$

This underpins the welfare finding. For every 10% increase in tariff-inclusive costs, foreign exporters slashed their base markups by 4%. By exercising terms-of-trade power via tariffs, the US successfully induced a squeeze on foreign profits, extracting rents as predicted by relevant trade theory (Ossa, 2014; Bagwell and Staiger, 1999), albeit at a large cost to aggregate efficiency. This analysis does not account for any foreign retaliation so we cannot fully outline the full cost to both US and global welfare.

Rationalizing the Markup Collapse A large fall in markups ($\phi_\mu < 0$) is required to generate incomplete base pass-through ($\tilde{\rho} \approx 0.6$) in the presence of strong scale economies. Standard theory suggests that as tariffs reduce volume, average costs rise, putting upward pressure on prices. For the base price to rise by less than the tariff, the firm must compress its markup enough to offset not only the tariff wedge but also the rising average costs.

This result, however, is not a modeling artifact but derived from three robust features of the data:

1. Significant Quantity Discounts ($\gamma \approx -0.3$).

[39] We use the identity $\tilde{\rho} = \rho(1 + \gamma\epsilon_D)$, substituting $\gamma_c \approx -0.28$.

2. Volume Contraction ($dQ/dT < 0$).

3. Incomplete Base Pass-Through ($\tilde{\rho} < 1$).

Any model that fixes markups (e.g., constant elasticity of demand) would counter-factually predict pass-through greater than 100%. Therefore, the data rejects constant-markup models in favor of a market structure with high strategic substitutability or demand curvature (Atkeson and Burstein, 2008; Gopinath and Itskhoki, 2011), where the markup adjustment is the important margin of adjustment.

Identified Market Structure These estimates rule out standard theoretical benchmarks. Following Weyl and Fabinger (2013), the pass-through rate is structurally determined by $\rho = (\mathcal{M} + \frac{\epsilon_D}{\epsilon_S})^{-1}$, where $\mathcal{M} \equiv \frac{\epsilon_D}{\epsilon_D - \theta}(1 + \theta\epsilon_{ms} + \theta\epsilon_\theta)$ is the structural markup term. Theoretically, CES demand with constant conduct (like monopolistic competition) implies a structural markup term of unity ($\mathcal{M} = 1$). However, imposing our observed pass-through ($\rho = 1$) and scale estimates returns a value of $\mathcal{M} \approx 1.4$. This deviation rejects the standard Melitz-Chaney benchmark, confirming that strategic complementarities, specifically variable markups, are quantitatively large. Under reasonable assumptions (e.g., linear demand (Melitz and Ottaviano, 2008), production function based-margins (Yue and Lin, 2023)), this implies a market structure of moderate oligopolistic competition ($\theta \in [0.2, 0.4]$), far removed from both perfect competition and the constant-markup behavior often assumed in trade models.

8 Conclusion

This paper shows the empirical relevance of quantity discounts at the transaction level for US imports and their importance for aggregate outcomes. These discounts appear to largely reflect scale economies on the supply-side and not market power. Furthermore, accounting for quantity discounts changes our understanding of 2018-2019 tariff pass-through and incidence.

Our decomposition shows that foreign exporters are effectively fighting a two-front war against the tariff. On one front, the tariff reduces trade volumes, pushing them up their average cost curves (the loss of scale economies). Our evidence suggests these costs are logistical—the rising per-unit cost of shipping, handling, and processing smaller batches. This "supply chain waste" represents a real efficiency loss, consuming resources that could have otherwise been deployed elsewhere. On the other front, to maintain market share, exporters compress their markups. Consequently, the incidence falls more

heavily on foreign producers than previously recognized. This finding challenges the view that US tariffs were largely paid by domestic consumers.

Going beyond this paper, the appropriate consideration of quantity discounts matters in a number of other contexts. For example, in any setting that uses material usage as an undistorted input, for example to measure distortions in output and labor markets (Yeh, Macaluso, and Hershbein, 2022). In industry studies, the scale of transactions may also lead to scale in firms, as in Ganapati (2025), complicating studies of sectoral welfare. In a macroeconomic context, such quantity discounts matter for papers that use aggregate input-output tables, or rely on a game where upstream seller prices are taken as a given by downstream firms (Baqaee and Farhi, 2020).

References

ADACHI, T. AND M. FABINGER (2022): "Pass-through, welfare, and incidence under imperfect competition," *Journal of Public Economics*, 211, 104589.

ALESSANDRIA, G., J. P. KABOSKI, AND V. MIDRIGAN (2010): "Inventories, lumpy trade, and large devaluations," *American Economic Review*, 100, 2304–2339.

ALLEN, T. (2014): "Information frictions in trade," *Econometrica*, 82, 2041–2083.

ALVIAREZ, V. I., M. FIORETTI, K. KIKKAWA, AND M. MORLACCO (2023): "Two-sided market power in firm-to-firm trade," Tech. rep., National Bureau of Economic Research.

AMITI, M., O. ITSKHOKI, AND J. KONINGS (2014): "Importers, exporters, and exchange rate disconnect," *American Economic Review*, 104, 1942–1978.

AMITI, M., S. J. REDDING, AND D. E. WEINSTEIN (2019): "The impact of the 2018 tariffs on prices and welfare," *Journal of Economic Perspectives*, 33, 187–210.

ANTRAS, P., T. C. FORT, AND F. TINTELNOT (2017): "The margins of global sourcing: Theory and evidence from US firms," *American Economic Review*, 107, 2514–2564.

ANTRAS, P. AND E. HELPMAN (2004): "Global sourcing," *Journal of political Economy*, 112, 552–580.

ANTWEILER, W. AND D. TREFLER (2002): "Increasing returns and all that: a view from trade," *American Economic Review*, 92, 93–119.

ARKOLAKIS, C., A. COSTINOT, D. DONALDSON, AND A. RODRÍGUEZ-CLARE (2019): "The elusive pro-competitive effects of trade," *The Review of Economic Studies*, 86, 46–80.

ARROW, K. J., T. HARRIS, AND J. MARSCHAK (1951): "Optimal inventory policy," *Econometrica: Journal of the Econometric Society*, 250–272.

ATALAY, E. (2014): "Materials prices and productivity," *Journal of the European Economic Association*, 12, 575–611.

ATKESON, A. AND A. BURSTEIN (2008): "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 98, 1998–2031.

BAGWELL, K. AND R. W. STAIGER (1999): "An economic theory of GATT," *American Economic Review*, 89, 215–248.

BAQAEE, D. R. AND E. FARHI (2020): "Productivity and misallocation in general equilibrium," *The Quarterly Journal of Economics*, 135, 105–163.

BARTELME, D. G., A. COSTINOT, D. DONALDSON, AND A. RODRIGUEZ-CLARE (2019): "The textbook case for industrial policy: Theory meets data," Tech. rep., National Bureau of Economic Research.

BASU, S. AND J. G. FERNALD (1997): "Returns to scale in US production: Estimates and implications," *Journal of political economy*, 105, 249–283.

BAUMOL, W. J. (1952): "The transactions demand for cash: An inventory theoretic approach," *The Quarterly journal of economics*, 66, 545–556.

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., E. DHYNE, G. MAGERMAN, K. MANOVA, AND A. MOXNES (2022): "The origins of firm heterogeneity: A production network approach," *Journal of Political Economy*, 130, 1765–1804.

BERRY, S. T. AND P. A. HAILE (2021): "Foundations of demand estimation," in *Handbook of industrial organization*, Elsevier, vol. 4, 1–62.

BISHOP, R. L. (1968): "The effects of specific and ad valorem taxes," *The Quarterly Journal of Economics*, 82, 198–218.

BORNSTEIN, G. AND A. PETER (2022): "Nonlinear pricing and misallocation," Tech. rep., Mimeo.

——— (2024): "Nonlinear pricing and misallocation," Tech. rep., National Bureau of Economic Research.

BRODA, C., N. LIMAO, AND D. E. WEINSTEIN (2008): "Optimal tariffs and market power: the evidence," *American Economic Review*, 98, 2032–2065.

BRODA, C. AND D. E. WEINSTEIN (2006): "Globalization and the Gains from Variety," *The Quarterly journal of economics*, 121, 541–585.

——— (2010): "Product creation and destruction: Evidence and price implications," *American Economic Review*, 100, 691–723.

BURSTEIN, A. T., J. CRAVINO, AND M. ROJAS (2024): *Input price dispersion across buyers and misallocation*, Banco Central de Chile.

CAVALLO, A., G. GOPINATH, B. NEIMAN, AND J. TANG (2021): "Tariff pass-through at the border and at the store: Evidence from us trade policy," *American Economic Review: Insights*, 3, 19–34.

CHEUNG, F. K. (1998): "Excise taxes on a non-uniform pricing monopoly: ad valorem and unit taxes compared," *Canadian Journal of Economics*, 1192–1203.

CHOR, D. AND B. LI (2024): "Illuminating the effects of the US-China tariff war on China's economy," *Journal of International Economics*, 150, 103926.

DELIPALLA, S. AND M. KEEN (1992): "The comparison between ad valorem and specific taxation under imperfect competition," *Journal of Public Economics*, 49, 351–367.

DUARTE, M., L. MAGNOLFI, D. QUINT, M. SØLVSTEN, AND C. SULLIVAN (2025): "Conduct and Scale Economies: Evaluating Tariffs in the US Automobile Market," *working paper*.

FAJGELBAUM, P., P. GOLDBERG, P. KENNEDY, A. KHANDELWAL, AND D. TAGLIONI (2024): "The US-China trade war and global reallocations," *American Economic Review: Insights*, 6, 295–312.

FAJGELBAUM, P. D., P. K. GOLDBERG, P. J. KENNEDY, AND A. K. KHANDELWAL (2020): "The return to protectionism," *The Quarterly Journal of Economics*, 135, 1–55.

FARROKHI, F. AND A. SODERBERY (2020): "Trade elasticities in general equilibrium," *Manuscript, Purdue University*.

FEENSTRA, R. C. (1994): "New product varieties and the measurement of international prices," *The American Economic Review*, 157–177.

FONTAINE, F., J. MARTIN, AND I. MEJEAN (2020): "Price discrimination within and across EMU markets: Evidence from French exporters," *Journal of international economics*, 124, 103300.

GANAPATI, S. (2025): "The modern wholesaler: Global sourcing, domestic distribution, and scale economies," *American Economic Journal: Microeconomics*, 17, 1–40.

GANAPATI, S. AND C. HOTTMAN (2025): "Micro and Macro Pass-Through: Tariffs and Taxes in a Structural Model of Supply and Demand," *Working Paper*.

GANAPATI, S., J. S. SHAPIRO, AND R. WALKER (2020): "Energy cost pass-through in US manufacturing: Estimates and implications for carbon taxes," *American Economic Journal: Applied Economics*, 12, 303–342.

GANAPATI, S., W. F. WONG, AND O. ZIV (2024): "Entrepot: Hubs, scale, and trade costs," *American Economic Journal: Macroeconomics*, 16, 239–278.

GERTLER, S. (2023): "The Structural Drivers of Price and Quantity Adjustment: Insights from Tariff and Exchange Rate Pass-through," Tech. rep., mimeo.

GOLDBERG, P. K. AND R. HELLERSTEIN (2008): "A structural approach to explaining incomplete exchange-rate pass-through and pricing-to-market," *American Economic Review*, 98, 423–429.

GOPINATH, G. AND O. ITSKHOKI (2011): "In search of real rigidities," *NBER Macroeconomics Annual*, 25, 261–310.

GOPINATH, G., O. ITSKHOKI, AND R. RIGOBON (2010): "Currency choice and exchange rate pass-through," *American Economic Review*, 100, 304–336.

HALPERN, L. AND M. KOREN (2007): "Pricing to firm: An analysis of firm-and product-level import prices," *Review of international economics*, 15, 574–591.

HANDLEY, K., F. KAMAL, AND R. MONARCH (2025): "Rising import tariffs, falling exports: When modern supply chains meet old-style protectionism," *American Economic Journal: Applied Economics*, 17, 208–238.

HILLBERRY, R. AND D. HUMMELS (2013): "Trade elasticity parameters for a computable general equilibrium model," in *Handbook of computable general equilibrium modeling*, Elsevier, vol. 1, 1213–1269.

HORNOK, C. AND M. KOREN (2015): "Per-shipment costs and the lumpiness of international trade," *Review of Economics and Statistics*, 97, 525–530.

HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 131, 1291–1364.

HUMMELS, D., V. LUGOVSKYY, AND A. SKIBA (2009): "The trade reducing effects of market power in international shipping," *Journal of Development Economics*, 89, 84–97.

HUMMELS, D. AND A. SKIBA (2004): "Shipping the good apples out? An empirical confirmation of the Alchian-Allen conjecture," *Journal of political Economy*, 112, 1384–1402.

IGNATENKO, A. (2020): "Price discrimination in international transportation: Evidence and implications," *mimeo*.

JENKIN, F. (1872): "3. on the principles which regulate the incidence of taxes," *Proceedings of the Royal Society of Edinburgh*, 7, 618–631.

KAMAL, F. AND R. MONARCH (2018): "Identifying foreign suppliers in US import data," *Review of International Economics*, 26, 117–139.

KAMAL, F. AND A. SUNDARAM (2016): "Buyer–seller relationships in international trade: Do your neighbors matter?" *Journal of International Economics*, 102, 128–140.

KIMBALL, M. S. (1995): "The quantitative analytics of the basic neomonetarist model," .

KLENOW, P. J. AND J. L. WILLIS (2016): "Real rigidities and nominal price changes," *Economica*, 83, 443–472.

KROLIKOWSKI, P. M. AND A. H. MCCALLUM (2025): "Tariffs and Goods-Market Search Frictions," Tech. rep., Federal Reserve.

KUGLER, M. AND E. VERHOOGEN (2012): "Prices, plant size, and product quality," *The Review of Economic Studies*, 79, 307–339.

LASHKARIPOUR, A. AND V. LUGOVSKYY (2023): "Profits, scale economies, and the gains from trade and industrial policy," *American Economic Review*, 113, 2759–2808.

MACKAY, A., N. H. MILLER, M. REMER, AND G. SHEU (2014): "Bias in reduced-form estimates of pass-through," *Economics Letters*, 123, 200–202.

MANOVA, K. AND Z. ZHANG (2012): "Export prices across firms and destinations," *The Quarterly Journal of Economics*, 127, 379–436.

MARSHALL, G. (2020): "Search and wholesale price discrimination," *The RAND Journal of Economics*, 51, 346–374.

MASKIN, E. AND J. RILEY (1984): "Monopoly with incomplete information," *The RAND Journal of Economics*, 15, 171–196.

MELESHCHUK, S. (2017): "Price Discrimination in International Trade: Empirical Evidence and Theory," Tech. rep., Tech. rep.

MELITZ, M. J. (2003): "The impact of trade on intra-industry reallocations and aggregate industry productivity," *econometrica*, 71, 1695–1725.

MELITZ, M. J. AND G. I. OTTAVIANO (2008): "Market size, trade, and productivity," *The review of economic studies*, 75, 295–316.

MRÁZOVÁ, M. AND J. P. NEARY (2017): "Not so demanding: Demand structure and firm behavior," *American Economic Review*, 107, 3835–3874.

MUNSON, C. L., J. JACKSON, ET AL. (2015): "Quantity discounts: An overview and practical guide for buyers and sellers," *Foundations and Trends® in Technology, Information and Operations Management*, 8, 1–130.

MUNSON, C. L. AND M. J. ROSENBLATT (1998): "Theories and realities of quantity discounts: An exploratory study," *Production and operations management*, 7, 352–369.

OSSA, R. (2014): "Trade wars and trade talks with data," *American Economic Review*, 104, 4104–4146.

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

SODERBERY, A. (2015): "Estimating import supply and demand elasticities: Analysis and implications," *Journal of International Economics*, 96, 1–17.

STIGLER, G. J. (1961): "The economics of information," *Journal of political economy*, 69, 213–225.

STOLE, L. A. (2007): "Price discrimination and competition," *Handbook of industrial organization*, 3, 2221–2299.

U.S. BUREAU OF LABOR STATISTICS (2025): "International Price Program (IPP): Data Sources," `https://www.bls.gov/opub/hom/ipp/data.htm`, accessed: 2025-10-24.

VERBOVEN, F. (2002): "Quality-based price discrimination and tax incidence: evidence from gasoline and diesel cars," *RAND Journal of Economics*, 275–297.

WEYL, E. G. AND M. FABINGER (2013): "Pass-through as an economic tool: Principles of incidence under imperfect competition," *Journal of political economy*, 121, 528–583.

YEH, C., C. MACALUSO, AND B. HERSHBEIN (2022): "Monopsony in the US labor market," *American Economic Review*, 112, 2099–2138.

YUE, W. AND Q. LIN (2023): "Export duration and firm markups: evidence from China," *Humanities and Social Sciences Communications*, 10, 1–11.

A Theoretical Derivations and Microfoundations

A.1 Derivation of Buyer Surplus Change

Buyer surplus is defined as the utility $U(q)$ derived from consuming quantity q minus the total expenditure. Given the non-linear unit price schedule $P(q) = \tilde{p}p(q)$, where $\tilde{p}$ is the tariff-inclusive base price level and $p(q)$ is the scale factor (reflecting quantity discounts), the buyer's total expenditure is $E(q) = q \cdot \tilde{p}p(q)$. The buyer chooses q to maximize surplus:

$$DS(\tilde{p}) = \max_{q} \{U(q) - \tilde{p}p(q)q\} \tag{24}$$

To determine the incidence of a tariff change, we first consider the effect of a change in the base price $\tilde{p}$. By the Envelope Theorem, the change in maximized surplus with respect to the parameter $\tilde{p}$ is equal to the partial derivative of the objective function, evaluating q at the optimal choice:

$$\frac{dDS}{d\tilde{p}} = -\frac{\partial E}{\partial \tilde{p}} = -p(q)q \tag{25}$$

This result mirrors Shephard's Lemma but is modified by the scale factor $p(q)$. It implies that the effective welfare loss from a base price increase is proportional to the *discounted* price paid, not the base price. Differentiating with respect to the tariff term $(1+\tau)$ and applying the chain rule yields:

$$\frac{dDS}{d(1+\tau)} = \int_t \frac{dDS_t}{d\tilde{p}_t} \frac{d\tilde{p}_t}{d(1+\tau)} dt = -\int_t q_t p(q_t) \frac{d\tilde{p}_t}{d(1+\tau)} dt \tag{26}$$

This corresponds to Equation (2) in the main text (which reports the magnitude of the incidence). If $p(q) = 1$ (uniform pricing), this collapses to the standard result where welfare loss is simply quantity times the change in price.

Simplification with Iso-elastic Pricing If we assume the scale factor takes the iso-elastic form $p(q) = q^{\gamma}$, then the term $p(q)q$ simplifies to $q^{\gamma} \cdot q = q^{1+\gamma}$. The buyer surplus derivative becomes:

$$\frac{dDS}{d(1+\tau)} = -\int_t q_t^{1+\gamma} \frac{d\tilde{p}_t}{d(1+\tau)} dt \tag{27}$$

This can be re-written in terms of observed transaction value. Let $V_t = P(q_t)q_t = \tilde{p}_t q_t^{1+\gamma}$ be the total value of transaction t. We can rewrite the derivative of the level term as

$\frac{d\tilde{p}_t}{d(1+\tau)} = \tilde{p}_t \frac{d \ln \tilde{p}_t}{d(1+\tau)}$. Substituting these in yields:

$$\frac{dDS}{d(1+\tau)} = -\int_t V_t \frac{d \ln \tilde{p}_t}{d(1+\tau)} dt \tag{28}$$

This result implies that if we define pass-through as the percentage change in the scale-free price ($\frac{d \ln \tilde{p}_t}{d \ln(1+\tau)}$), aggregate incidence is simply the pass-through multiplied by total transaction value. This is the expression used in the main text.

Caveat on Large Adjustments It is important to qualify that the Envelope Theorem result relies on the assumption of marginal adjustments. If the inventory or administrative cost function $C(q)$ is convex (e.g., rising sharply as order size approaches zero due to fixed logistical overhead), the linearization implied by the Envelope theorem will understate the welfare loss. The lost surplus from the 'intramarginal' units that are no longer purchased—and the lost efficiency from the units now purchased at inefficiently small scales—represents a second-order deadweight loss (analogous to a Harberger triangle, but potentially larger).

A.1.1 Impact on Incidence: γ is Endogenous

When the shape parameter γ is endogenous, the incidence of the tax becomes heterogeneous across consumers of different sizes. Applying the Envelope Theorem to the consumer's problem, the change in consumer surplus is driven by the direct shift in the price schedule at the chosen quantity q:

$$\frac{dCS(q)}{d \log T} = -\text{Expenditure}(q) \times (\tilde{\rho} + \rho_\gamma \log q) \tag{29}$$

This reveals that the "effective" pass-through for welfare analysis is not a single number but a function of quantity: $\tilde{\rho}_{\text{eff}}(q) = \tilde{\rho} + \rho_\gamma \log q$. Consequently, relying on the aggregate average $\tilde{\rho}$ may mask distributional effects where the tax burden is shifted disproportionately onto large or small buyers depending on the sign of ρ_γ.

A.2 Derivation of Producer Surplus Change

In this section, we derive the change in producer surplus with respect to tariffs as presented in Equation 3.

Producer surplus is defined as the integral of profits over all transactions:

$$PS = \int_{t \in T} \pi(t)\, dt \tag{30}$$

$$= \int_{t \in T} [R(t) - C(t)]\, dt \tag{31}$$

where $R(t)$ denotes revenue and $C(t)$ denotes total variable costs for transaction t. Note that revenue is a function of the tariff-inclusive price faced by the buyer, stripped of the tariff wedge. Let τ be the ad-valorem tariff rate.

The revenue for a single transaction is:

$$R(t) = \frac{P(q_t)}{1+\tau} \cdot q_t = \frac{\tilde{p} \cdot p(q_t)}{1+\tau} \cdot q_t \tag{32}$$

The total cost for a single transaction is:

$$C(t) = \text{Unit Cost}(q_t) \cdot q_t = [\tilde{c} \cdot c(q_t)] \cdot q_t \tag{33}$$

We are interested in the derivative of aggregate producer surplus with respect to the tariff factor $(1+\tau)$. By linearity of the integral, we can differentiate term by term:

$$\frac{dPS}{d(1+\tau)} = \int_t \left[\frac{dR(t)}{d(1+\tau)} - \frac{dC(t)}{d(1+\tau)} \right] dt \tag{34}$$

A.2.1 Revenue Decomposition

Differentiating the logarithm of revenue $R(t)$ with respect to $\log(1+\tau)$:

$$\log R = \log \tilde{p} + \log p(q) + \log q - \log(1+\tau) \tag{35}$$

$$\frac{d \log R}{d \log(1+\tau)} = \frac{d \log \tilde{p}}{d \log(1+\tau)} + \frac{d \log p(q)}{d \log q} \frac{d \log q}{d \log(1+\tau)} + \frac{d \log q}{d \log(1+\tau)} - 1 \tag{36}$$

Grouping terms involving q:

$$\frac{d \log R}{d \log(1+\tau)} = \frac{d \log \tilde{p}}{d \log(1+\tau)} + \left(1 + \frac{d \log p(q)}{d \log q}\right) \frac{d \log q}{d \log(1+\tau)} - 1 \tag{37}$$

Converting from elasticities back to levels using $dX/X = d \log X$:

$$\frac{dR}{d(1+\tau)} = \left[\frac{d \log \tilde{p}}{d \log(1+\tau)} + (1+\gamma) \frac{d \log q}{d \log(1+\tau)} - 1 \right] \frac{R}{1+\tau} \tag{38}$$

where $\gamma = \frac{d \log p(q)}{d \log q}$ is the elasticity of the price schedule (quantity discount).

A.2.2 Cost Decomposition

Similarly, for costs $C(t)$:

$$\log C = \log \tilde{c} + \log c(q) + \log q \tag{39}$$

$$\frac{d \log C}{d \log(1+\tau)} = \frac{d \log \tilde{c}}{d \log(1+\tau)} + \left(1 + \frac{d \log c(q)}{d \log q}\right) \frac{d \log q}{d \log(1+\tau)} \tag{40}$$

Assuming base costs $\tilde{c}$ are constant with respect to tariffs (no input linkages): $\frac{d \log \tilde{c}}{d \log(1+\tau)} = 0$. However, retaining the general form:

$$\frac{dC}{d(1+\tau)} = \left[\frac{d \log \tilde{c}}{d \log(1+\tau)} + (1+\gamma^c) \frac{d \log q}{d \log(1+\tau)}\right] \frac{C}{1+\tau} \tag{41}$$

Substituting these components back into the integral yields the result in the main text.

A.3 Incidence - Ex-ante supply side knowledge

The above derivation did not use an envelope condition for the producer and did not assume knowledge of market structure. This has both positives and negatives. In particular, why we may need to know the level of markups, we do not need to know the level from which firms internalize their pricing decisions.

We determine the incidence I, defined as the ratio of Consumer Surplus loss to Producer Surplus loss (relative to initial spend/revenue) for a small tax increase.

A.3.1 Consumer Surplus

The consumer maximizes $CS = \max_q [U(q) - \tilde{p} q^{1+\gamma}]$. By the Envelope Theorem, the change in welfare ignores the second-order effect of re-optimizing q. The loss is driven entirely by the direct effect of the parameter change:

$$\frac{dCS}{dT} = \frac{\partial \mathcal{L}}{\partial \tilde{p}} \frac{d\tilde{p}}{dT} = -q^{1+\gamma} \frac{d\tilde{p}}{dT} \tag{42}$$

Expressed as an elasticity with respect to the tax T:

$$\frac{dCS}{d \log T} = -(\tilde{p} q^{1+\gamma}) \frac{d \log \tilde{p}}{d \log T} = -\text{Expenditure} \times \tilde{\rho} \tag{43}$$

A.3.2 Producer Surplus: Structural Approach

Assuming an ad-valorem tax collected from revenue, industry profit is $\Pi = \frac{p(Q)Q}{T} - NC(Q/N)$. For $\theta < 1$, firms do not jointly maximize industry profit; they produce excessive output relative to the cartel optimum. A tax increase induces an output contraction, which has a positive strategic effect on firm profits that partially offsets the direct tax burden.

$$\frac{dPS}{d\log T} \approx -\text{Revenue}\,[1-(1-\theta)\phi] \tag{44}$$

A.3.3 Relative Incidence

The incidence ratio $I = \frac{dCS/d\log T}{dPS/d\log T}$ is:

$$I = \frac{-\text{Expenditure} \times \tilde{\rho}}{-\text{Revenue} \times [1-(1-\theta)\phi]} \approx \frac{\tilde{\rho}}{1-(1-\theta)\phi} \tag{45}$$

A.3.4 Interpretation

- **Monopoly ($\theta = 1$):** The strategic effect vanishes (the bracketed term becomes 1). $I = \tilde{\rho}$.
- **Competition ($\theta \to 0$):** The denominator approaches $1 - \phi$.
- **Role of $\tilde{\rho}$ vs ρ:**
  - The numerator uses $\tilde{\rho}$ because consumer welfare loss is driven by the *shift in the price schedule* (Envelope Theorem).
  - The denominator uses ϕ because the firm's strategic reaction is driven by the *contraction in volume*, which depends on the realized unit price consumers face.

A.4 Recovering Pass-Through

We follow Weyl and Fabinger (2013), but with both ad-valorem taxes and endogenous pricing (subject to caveats).

A.4.1 Demand and Pricing

Consumers face a non-linear price schedule set by the firm. The unit price p depends on quantity q according to the schedule:

$$p(q) = \tilde{p} \cdot q^{\gamma} \tag{46}$$

where γ is a fixed parameter capturing the non-linearity (e.g., bulk discounts if $\gamma < 0$, surcharges if $\gamma > 0$).

Assumption (Consumer Internalization): The consumer fully observes and internalizes the shape of this schedule (the quantity discount or surcharge determined by γ), but takes the base price parameter $\tilde{p}$ as given. They maximize utility $U(q) - p(q)q$. The First Order Condition ($U'(q) = p(q) + p'(q)q$) yields the inverse demand curve effectively faced by the firm:

$$p(q) = \frac{U'(q)}{1+\gamma} \tag{47}$$

A.4.2 Firm Behavior

The firm chooses the base price $\tilde{p}$ to maximize profit, given the schedule shape γ and a conduct parameter θ (where $\theta = 0$ implies perfect competition and $\theta = 1$ implies monopoly). The firm faces an ad-valorem tax (or tariff) $T = 1 + \tau$. The equilibrium condition is:

$$MR(q) = MC(q) \cdot T \tag{48}$$

Marginal Revenue is defined structurally by conduct θ:

$$MR(q) = p\left(1 - \frac{\theta}{\epsilon_D}\right) \tag{49}$$

A.4.3 Elasticity Definitions

To derive pass-through, we define the following structural elasticities:

- **Demand Elasticity (ϵ_D):** Defined on the inverse demand curve $p(q)$ faced by the firm:
$$\frac{1}{\epsilon_D} \equiv -\frac{dp}{dq}\frac{q}{p}$$
- **Marginal Surplus (ms):** Defined as the absolute gap between price and marginal revenue for a monopolist (or the slope of the inverse demand times quantity). It

represents the consumer's inframarginal surplus on the last unit:

$$ms(q) \equiv -p'(q)q = \frac{p(q)}{\epsilon_D}$$

Note that we can rewrite Marginal Revenue as $MR(q) = p(q) - \theta ms(q)$.

- **Curvature of Marginal Surplus (ϵ_{ms}^{-1}):** The elasticity of the marginal surplus function, which captures the curvature of demand (often related to the convexity of the demand curve):

$$\epsilon_{ms}^{-1} \equiv \frac{dms}{dq}\frac{q}{ms}$$

- **Conduct Elasticity (ϵ_{θ}^{-1}):** While θ is often treated as a parameter, it may vary with quantity (e.g., if collusion stability depends on volume). We define its elasticity as:

$$\epsilon_{\theta}^{-1} \equiv \frac{d\theta}{dq}\frac{q}{\theta}$$

A.4.4 Derivation of Unit Price Pass-Through (ρ)

We seek the pass-through rate of the tax to the observable unit price, defined as $\rho \equiv \frac{d\log p}{d\log T}$.

A.4.5 Log-Linearization

Taking logs of the equilibrium condition $MR = MC \cdot T$:

$$\log MR(q) = \log MC(q) + \log T \tag{50}$$

Differentiating with respect to $\log q$ and multiplying by the response of quantity to tax $\frac{d\log q}{d\log T}$:

$$\frac{d\log MR}{d\log q}\frac{d\log q}{d\log T} = \frac{d\log MC}{d\log q}\frac{d\log q}{d\log T} + 1 \tag{51}$$

Rearranging to solve for the quantity response:

$$\frac{d\log q}{d\log T} = \frac{1}{\frac{d\log MR}{d\log q} - \frac{d\log MC}{d\log q}} \tag{52}$$

By definition of demand elasticity, the price change is linked to quantity by $\frac{d\log p}{d\log q} = -\frac{1}{\epsilon_D}$. Thus $\rho = -\frac{1}{\epsilon_D}\frac{d\log q}{d\log T}$. Substituting this in:

$$\rho = \frac{-1/\epsilon_D}{\frac{d\log MR}{d\log q} - \frac{1}{\epsilon_S}} = \frac{1}{-\epsilon_D \frac{d\log MR}{d\log q} + \frac{\epsilon_D}{\epsilon_S}} \tag{53}$$

A.4.6 Expansion of Marginal Revenue

We expand the term $\frac{d\log MR}{d\log q}$. Using $MR = p - \theta ms$ and differentiating with respect to q:

$$\begin{aligned}\frac{dMR}{dq} &= p' - (\theta' ms + \theta ms') \\ &= p' - \theta ms\left(\frac{\theta'}{\theta} + \frac{ms'}{ms}\right)\end{aligned}$$

Converting to elasticities. Recall $p' = -p/(q\epsilon_D)$ and $ms = p/\epsilon_D$:

$$\frac{dMR}{dq} = \frac{dp}{dq}\left[1 + \theta\frac{ms}{-p'}\left(\frac{\epsilon_\theta}{q} + \frac{\epsilon_{ms}}{q}\right)\right] \tag{54}$$

Note that $\frac{ms}{-p'} = \frac{-p'q}{-p'} = q$. Thus:

$$\frac{dMR}{dq} = \frac{dp}{dq}\left[1 + \theta(\epsilon_\theta^{-1} + \epsilon_{ms}^{-1})\right] \tag{55}$$

Finally, converting to the logarithmic derivative $\frac{d\log MR}{d\log q} = \frac{dMR}{dq}\frac{q}{MR}$:

$$\frac{d\log MR}{d\log q} = \underbrace{\left(\frac{dp}{dq}\frac{q}{p}\right)}_{-\frac{1}{\epsilon_D}}\frac{p}{MR}[1 + \theta\epsilon_{ms}^{-1} + \theta\epsilon_\theta^{-1}] \tag{56}$$

A.4.7 Pass-Through

Substituting this back into Eq (53), and using the identity $p/MR = \epsilon_D/(\epsilon_D - \theta)$, we substitute $\frac{p}{MR}$ and factor it out to obtain the structural pass-through formula:

$$\rho = \frac{1}{\frac{\epsilon_D}{\epsilon_D - \theta}\left[1 + \theta\epsilon_{ms}^{-1} + \theta\epsilon_\theta^{-1} + \frac{\epsilon_D - \theta}{\epsilon_S}\right]} \tag{57}$$

Note: This factorization highlights that pass-through is scaled by the inverse markup. The formula depends on γ only implicitly (as γ affects the equilibrium values of ϵ_D and ϵ_S), but the structural form is identical to the standard case.

Comparison to the Specific-Tax Formulation This is a scaled version of the specific tax formulation in Weyl and Fabinger (2013). In particular:

$$\rho = \rho_{\text{ST}} \cdot \frac{mr}{p} = \rho_{\text{ST}} \cdot \frac{\epsilon_D - \theta}{\epsilon_D}.$$

Where the ad-valorem pass-through $\rho = \frac{d \log p}{d \log T}$ is linked to the specific tax pass-through $\rho_{\text{ST}} = \frac{dp}{dt}$ via an "equivalent" specific tax shock.

There are a few differences. The elasticity formulation means we do not use the Lerner markup rule, but rather just the internalized marginal revenue term and do not worry about the specific tax starting at $t = 0$. The relationship between specific tax t and ad-valorem tax $1 + \tau$ is exactly inverse of the markup.

Let $t(T)$ be a specific tax that yields the same equilibrium marginal cost shift as the ad-valorem tax T at the current optimum:

$$t(T) = (T - 1)MC$$

Differentiating this effective burden with respect to T gives the intensity of the cost shock:

$$\frac{dt}{dT} = MC$$

Using the chain rule, the price response to the ad-valorem tax is the specific pass-through rate times this cost shock intensity:

$$\frac{dp}{dT} = \frac{dp}{dt}\frac{dt}{dT} = \rho_{\text{ST}} \cdot MC$$

To convert this to an elasticity, we multiply by $\frac{T}{p}$:

$$\rho = \frac{T}{p}\frac{dp}{dT} = \frac{T}{p}(\rho_{\text{ST}} \cdot MC) = \rho_{\text{ST}}\frac{MC \cdot T}{p}$$

Using the equilibrium condition $MR = MC \cdot T$, we substitute the numerator:

$$\rho = \rho_{\text{ST}}\frac{MR}{p}$$

A.4.8 Base Price Pass-Through ($\tilde{\rho}$)

While ρ measures the change in the realized unit price, welfare analysis requires tracking the shift in the pricing *schedule*. We define the base price pass-through as:

$$\tilde{\rho} \equiv \frac{d \log \tilde{p}}{d \log T} \tag{58}$$

Relationship between ρ and $\tilde{\rho}$ From the pricing schedule $p = \tilde{p} q^{\gamma}$, we differentiate with respect to $\log T$:

$$\rho = \tilde{\rho} + \gamma \frac{d \log q}{d \log T} = \tilde{\rho} + \gamma(-\epsilon_D \rho) \tag{59}$$

Solving for $\tilde{\rho}$:

$$\tilde{\rho} = \rho(1 + \gamma \epsilon_D) \tag{60}$$

This geometric link holds regardless of market structure or conduct. It relates the movement *along* the schedule (determined by ρ) to the shift *of* the schedule ($\tilde{\rho}$).

A.5 Recovery of Pass-Through from Markup Elasticities

We can alternatively recover ρ using the elasticities of the cost and markup components.

Assumption (Decomposition): We decompose the pricing schedule $p(q)$ into a cost schedule $c(q)$ and a markup schedule $\mu(q)$ such that $p(q) = c(q)\mu(q)T$. We further assume the base parameters decompose similarly:

$$\tilde{p} = \tilde{c} \cdot \tilde{\mu} \cdot T \tag{61}$$

where $\tilde{c}$ is the base cost and $\tilde{\mu}$ is the base markup.

We define the local elasticities (log-slopes) and the tax-sensitivities (base pass-throughs):

- $\gamma = \gamma_c + \gamma_\mu$ (Schedule slopes: Price = Cost + Markup)
- $\phi_c \equiv \frac{d \log \tilde{c}}{d \log T}$ (Base cost pass-through)
- $\phi_\mu \equiv \frac{d \log \tilde{\mu}}{d \log T}$ (Base markup pass-through)

From the decomposition of $\tilde{p}$, the total shift in the base price is the sum of the shifts in its components:

$$\tilde{\rho} = \frac{d \log \tilde{p}}{d \log T} = \phi_c + \phi_\mu + 1 \tag{62}$$

Substituting this into the geometric link $\tilde{\rho} = \rho(1 + \gamma\epsilon_D)$ from Section 3:

$$\rho(1 + \gamma\epsilon_D) = 1 + \phi_\mu + \phi_c \tag{63}$$

Solving for ρ:

$$\rho = \frac{1 + \phi_\mu + \phi_c}{1 + \gamma\epsilon_D} \tag{64}$$

Relation to Structural Formula: Comparing Equation (64) with the original structural pass-through formula, we can establish a mapping between the observable markup elasticities and the underlying structural parameters. The denominator $1 + \gamma\epsilon_D$ corresponds to the structural denominator, implying that the standard "curvature" terms (ϵ_{ms}) are implicitly captured by γ_μ. Importantly, the numerator term ϕ_μ is generally non-zero in the Weyl-Fabinger framework. As the tax induces a quantity contraction, the firm moves to a new point on the demand curve where the optimal markup is different; to maintain the fixed functional form of the price schedule ($p = \tilde{p}q^\gamma$), the base markup $\tilde{p}$ must adjust. Thus, ϕ_μ acts as a sufficient statistic for the endogenous reaction of the markup to the tax-induced volume loss.

Algebraic Identity for ϕ_μ: We can explicitly express ϕ_μ by using the identity for the base price pass-through derived above: $\tilde{\rho} = 1 + \phi_\mu + \phi_c$. Assuming base costs are unresponsive to the tax ($\phi_c = 0$) and rearranging, we find a simple identity relating the base markup response to the observable base price pass-through:

$$\phi_\mu = \tilde{\rho} - 1 \tag{65}$$

This indicates that the elasticity of the base markup (ϕ_μ) captures exactly the extent to which the base price pass-through ($\tilde{\rho}$) deviates from unity. If $\tilde{\rho} < 1$ (base price rises by less than the tax), it implies $\phi_\mu < 0$ (the firm compresses its base markup). Conversely, if $\tilde{\rho} > 1$, the firm expands its base markup.

A.5.1 Endogenous Shape Parameter γ

If the seller can adjust the non-linearity parameter γ in response to the tax, the pass-through rate becomes heterogeneous across buyers of different sizes. Let the price schedule be $p(q, T) = \tilde{p}(T)q^{\gamma(T)}$. Differentiating $\log p$ with respect to $\log T$:

$$\rho(q) = \tilde{\rho} + \rho_\gamma \log q + \gamma \frac{d \log q}{d \log T} \tag{66}$$

where $\rho_\gamma \equiv \frac{d\gamma}{d\log T}$ captures the steepening or flattening of the schedule.

The consumer chooses quantity q where $p(q) = \frac{U'(q)}{1+\gamma}$. Taking logs and differentiating with respect to $\log T$:

$$\rho(q) = \frac{d\log U'}{d\log q}\frac{d\log q}{d\log T} - \frac{d\log(1+\gamma)}{d\gamma}\rho_\gamma \tag{67}$$

Using the identity $\epsilon_D^{-1} = -\frac{d\log U'}{d\log q}$ (assuming $p \approx U'$ for elasticity scaling), we have:

$$\rho(q) = -\frac{1}{\epsilon_D}\frac{d\log q}{d\log T} - \frac{\rho_\gamma}{1+\gamma} \tag{68}$$

This implies that the quantity response is driven by both the price change and the shift in the effective demand curve due to γ:

$$\frac{d\log q}{d\log T} = -\epsilon_D\left(\rho(q) + \frac{\rho_\gamma}{1+\gamma}\right) \tag{69}$$

Substituting this quantity response back into the price decomposition:

$$\rho(q) = \tilde{\rho} + \rho_\gamma \log q - \gamma\epsilon_D\left(\rho(q) + \frac{\rho_\gamma}{1+\gamma}\right) \tag{70}$$

Grouping terms by $\rho(q)$ yields the generalized pass-through formula:

$$\rho(q) = \frac{\tilde{\rho} + \rho_\gamma\left(\log q - \frac{\gamma\epsilon_D}{1+\gamma}\right)}{1+\gamma\epsilon_D} \tag{71}$$

Implications:

- **Size-Dependence:** If $\rho_\gamma \neq 0$, pass-through depends on buyer size ($\log q$). If the firm steepens the curve ($\rho_\gamma > 0$), larger buyers face higher pass-through.
- **Demand Dampening:** The term $-\frac{\gamma\epsilon_D}{1+\gamma}$ reflects that an increase in γ acts like a tax on the marginal unit, reducing quantity demanded and thus dampening the equilibrium price increase needed.

A.6 Generalizing the Framework: The Extensive Margin

Our local incidence analysis in Section 2.2 captures the intensive margin (changes in transaction size) but omits the extensive margin (changes in transaction frequency). We now

extend the framework to incorporate this margin and provide a complete welfare calculation.

The change in producer surplus, $dPS = dR - dC$, can be calculated using our empirical findings. Since quantity discounts are cost-driven ($\gamma = \gamma_c$) and scale-free costs are unaffected by tariffs ($d\ln\tilde{c}/d\ln(1+\tau) = 0$), we can relate costs to revenues via the markup: $C = R/\tilde{\mu}$. The change in producer surplus is then:

$$\frac{dPS}{d(1+\tau)} = \frac{R}{1+\tau}\left[\frac{d\ln R}{d\ln(1+\tau)}\left(1-\frac{1}{\tilde{\mu}}\right) + \frac{1}{\tilde{\mu}}\frac{d\ln\tilde{\mu}}{d\ln(1+\tau)}\right]. \tag{72}$$

All terms on the right-hand side are either directly estimated from our tariff regressions or taken from the literature, allowing us to compute the incidence on producers without direct cost data.

The downstream buyer surplus calculation in Equation (2) is incomplete, as it only captures the welfare change for continuing transactions (the intensive margin). Our results show a strong extensive margin response, with a significant drop in the number of transactions. A more complete expression for the change in downstream buyer surplus is:

$$\frac{dDS}{d(1+\tau)} = \int_{t\in\mathcal{T}_{cont}} \frac{d\tilde{p}_t}{d(1+\tau)} q_t dt - \int_{t\in\mathcal{T}_{exit}} DS_t dt, \tag{73}$$

where $\mathcal{T}_{cont}$ is the set of continuing transactions, $\mathcal{T}_{exit}$ is the set of exiting transactions, and DS_t is the surplus from a transaction prior to its exit. While precisely calculating the lost surplus from exiting transactions (DS_t) is difficult without a full demand model.

A simple approximation is to assume that the average surplus lost per exiting transaction is proportional to the average surplus lost per continuing transaction. If we further assume that the pre-tariff quantity of exiting transactions is similar to that of continuing ones, we can approximate the total lost surplus from the extensive margin by scaling the intensive margin loss by the relative number of exiting to continuing transactions. This leads to the following approximation for the total change in downstream buyer surplus:

$$\frac{dDS}{d(1+\tau)} = \frac{d(T\cdot\overline{DS})}{d(1+\tau)} = \frac{dT}{d(1+\tau)}\overline{DS} + T\frac{d\overline{DS}}{d(1+\tau)}, \tag{74}$$

where T is the number of transactions and $\overline{DS}$ is the average downstream buyer surplus per transaction. The first term captures the extensive margin (change in the number of transactions), while the second term captures the intensive margin (change in surplus per transaction). We can approximate the change in average surplus using our intensive margin estimate for continuing transactions: $T\frac{d\overline{DS}}{d(1+\tau)} \approx \int_{t\in\mathcal{T}_{cont}} \frac{d\tilde{p}_t}{d(1+\tau)} q_t dt$.

A.7 Model of Inventories

A.7.1 Baseline: Linear Pricing with Fixed Costs

Consider a firm that faces a constant downstream demand rate D per period. The firm purchases inputs in batches of size Q. The unit price of the input is fixed at p. However, the firm incurs a fixed ordering cost A for every transaction placed. Additionally, holding inventory incurs a cost. Let h denote the holding cost per unit per period (which may be interpreted as the interest rate times the unit price, $i \cdot p$).

The firm chooses the order size Q to minimize the total cost per period, which consists of purchasing costs, ordering costs, and inventory holding costs. The number of orders per period is D/Q. Assuming inventory is depleted linearly, the average inventory level is $Q/2$.

The total cost function is:

$$TC(Q) = \underbrace{Dp}_{\text{Purchase Cost}} + \underbrace{\frac{D}{Q}A}_{\text{Ordering Cost}} + \underbrace{\frac{Q}{2}h}_{\text{Holding Cost}} \tag{75}$$

Minimizing with respect to Q, the first-order condition is:

$$\begin{aligned} -\frac{D}{Q^2}A + \frac{h}{2} &= 0 \\ \frac{h}{2} &= \frac{DA}{Q^2} \end{aligned}$$

Solving for the optimal order quantity Q^*:

$$Q^* = \sqrt{\frac{2AD}{h}} \tag{76}$$

This is the standard Economic Order Quantity (EOQ) result. We can express the relationship in logarithms to find the elasticity of order size with respect to demand:

$$\ln Q^* = \frac{1}{2}\left(\ln 2 + \ln A - \ln h\right) + \frac{1}{2}\ln D \tag{77}$$

Thus, the elasticity of order size with respect to demand is:

$$\frac{d \ln Q^*}{d \ln D} = \frac{1}{2} \tag{78}$$

Similarly, the number of transactions $T = D/Q$ scales with demand as:

$$\ln T = \ln D - \ln Q^* = \frac{1}{2} \ln D + \text{constant} \tag{79}$$

implying an elasticity of transactions with respect to demand of 0.5.

A.7.2 Extension: Non-Linear Pricing

We now generalize the model to allow for non-linear pricing, specifically quantity discounts. Suppose there are no fixed ordering costs ($A = 0$), but the unit price $p(q)$ depends on the order size q. We assume an iso-elastic price schedule:

$$p(q) = \bar{p} q^{\gamma} \tag{80}$$

where $\gamma < 0$ represents the elasticity of the unit price with respect to order size. A negative γ implies a quantity discount.

The firm's problem is to minimize the sum of purchasing and holding costs:

$$\min_{q} \quad D p(q) + \frac{q}{2} h \tag{81}$$

Substituting the functional form for price:

$$\min_{q} \quad D \bar{p} q^{\gamma} + \frac{q}{2} h \tag{82}$$

The first-order condition is:

$$\begin{aligned} D \bar{p} \gamma q^{\gamma - 1} + \frac{h}{2} &= 0 \\ -D \bar{p} \gamma q^{\gamma - 1} &= \frac{h}{2} \end{aligned}$$

Solving for q:

$$\begin{aligned} q^{\gamma - 1} &= -\frac{h}{2D\bar{p}\gamma} \\ q &= \left(-\frac{h}{2D\bar{p}\gamma} \right)^{\frac{1}{\gamma - 1}} \end{aligned}$$

Taking logs to determine the elasticity:

$$\ln q = \frac{1}{\gamma - 1}\left[\ln\left(-\frac{h}{2\bar{p}\gamma}\right) - \ln D\right]$$
$$\ln q = \text{constant} - \frac{1}{\gamma - 1}\ln D$$
$$\ln q = \text{constant} + \frac{1}{1 - \gamma}\ln D$$

The elasticity of order size with respect to demand is:

$$\frac{d\ln q}{d\ln D} = \frac{1}{1-\gamma} \tag{83}$$

If the price elasticity is $\gamma \approx -1/3$, as suggested by the data, the order size elasticity becomes:

$$\frac{1}{1-(-1/3)} = \frac{1}{4/3} = 0.75 \tag{84}$$

This model predicts a higher pass-through of demand shocks to order size compared to the fixed-cost model (0.75 vs 0.50), as the firm is incentivized to increase order sizes not just to amortize fixed costs, but also to access lower unit prices.

A.7.3 Comparative Statics: The Impact of Tariffs

We now consider the impact of an ad-valorem tariff τ on the optimal order size and order frequency, explicitly allowing for incomplete pass-through to the base price.

Let $p(\tau)$ be the exporter's base unit price. The landed unit cost faced by the buyer is $P(\tau) = p(\tau)(1+\tau)$. We define the pass-through elasticity of the landed price with respect to the tariff factor as ρ:

$$\rho \equiv \frac{d\ln P(\tau)}{d\ln(1+\tau)} = 1 + \frac{d\ln p(\tau)}{d\ln(1+\tau)} \tag{85}$$

If pass-through is complete, $\rho = 1$ (exporter price p is constant). If exporters lower markups, $\rho < 1$.

We assume inventory holding costs are proportional to the landed value of the stock, so the holding cost per unit is $h(\tau) = i \cdot P(\tau)$.

Case 1: Fixed Cost Model In the baseline model with fixed ordering costs A, the total cost function is:

$$TC(Q, \tau) = DP(\tau) + \frac{D}{Q}A + \frac{Q}{2}iP(\tau) \tag{86}$$

The optimal order quantity Q^* is determined by balancing ordering and holding costs:

$$Q^*(\tau) = \sqrt{\frac{2DA}{iP(\tau)}} \tag{87}$$

Taking logs and differentiating with respect to $\ln(1+\tau)$:

$$\ln Q^*(\tau) = \frac{1}{2}\ln(2DA) - \frac{1}{2}\ln i - \frac{1}{2}\ln P(\tau) \tag{88}$$

$$\frac{d\ln Q^*}{d\ln(1+\tau)} = -\frac{1}{2}\frac{d\ln P(\tau)}{d\ln(1+\tau)} = -\frac{1}{2}\rho \tag{89}$$

The model predicts that order size decreases with tariffs, provided that pass-through is positive ($\rho > 0$). Importantly, the magnitude of this quantity reduction depends on the pass-through rate. If exporters absorb some tariff cost ($\rho < 1$), the reduction in order size is dampened but not eliminated.

Similarly, the order frequency $T = D/Q$ increases:

$$\frac{d\ln T^*}{d\ln(1+\tau)} = \frac{1}{2}\rho \tag{90}$$

Case 2: Non-Linear Pricing Model Consider the model with non-linear pricing $p(q) = \bar{p}(\tau)q^\gamma$ and no fixed ordering costs. The landed unit price schedule is $P(q, \tau) = \bar{p}(\tau)q^\gamma(1+\tau)$. Let $\bar{P}(\tau) = \bar{p}(\tau)(1+\tau)$ denote the level shifter of the landed price schedule.

The total cost function is:

$$TC(q, \tau) = D\bar{P}(\tau)q^\gamma + \frac{q}{2}i\bar{P}(\tau)q^\gamma \tag{91}$$

Notice that the price level term $\bar{P}(\tau)$ factors out of the expression:

$$TC(q, \tau) = \bar{P}(\tau)\left[Dq^\gamma + \frac{i}{2}q^{\gamma+1}\right] \tag{92}$$

The optimal q^* minimizes the term in the brackets. Since $\bar{P}(\tau)$ is merely a multiplicative scaling factor (assuming $\bar{P}(\tau) > 0$), it does not affect the first-order condition with respect

to q:

$$\frac{d}{dq}\left[Dq^{\gamma} + \frac{i}{2}q^{\gamma+1}\right] = D\gamma q^{\gamma-1} + \frac{i}{2}(\gamma+1)q^{\gamma} = 0 \tag{93}$$

Solving for q^*:

$$q^* = \frac{-2D\gamma}{i(\gamma+1)} \tag{94}$$

Crucially, q^* depends only on demand D, the interest rate i, and the curvature γ. It is independent of the price level $\bar{P}(\tau)$.

$$\frac{dq^*}{d\tau} = 0 \tag{95}$$

In this pure non-linear pricing framework, the optimal order size is invariant to the tariff, regardless of the pass-through rate ρ. The tariff scales both the marginal benefit of buying (the unit price) and the marginal cost of holding (the inventory value) by the exact same factor ρ, leaving the trade-off unchanged.

Relation to Empirical Results The empirical analysis in Section 7 documents two key facts:

1. **Quantity Response:** Transaction sizes fall significantly in response to tariffs ($\frac{d\ln q}{d\ln(1+\tau)} \approx -0.6$).

2. **Price Response:** Pass-through to the base price schedule is incomplete ($\rho < 1$), meaning exporters lower $\bar{p}(\tau)$.

The significant negative quantity response rejects the pure Non-Linear Pricing model (Case 2) and supports the Fixed Cost model (Case 1). The fact that quantities fall despite incomplete pass-through ($\rho < 1$) further validates the robustness of the fixed-cost mechanism: even a partial increase in landed holding costs is sufficient to trigger logistical adjustments.

This mechanism has a feedback effect on observed prices. As firms reduce Q to save on holding costs, they mechanically move up the unit price schedule $p(q) \propto q^{\gamma}$. The observed change in unit price is therefore:

$$\underbrace{\frac{d\ln P_{obs}}{d\ln(1+\tau)}}_{\text{Observed Pass-Through}} = \underbrace{\rho}_{\text{Shift in Schedule}} + \underbrace{\gamma\frac{d\ln q}{d\ln(1+\tau)}}_{\text{Mechanical Scale Effect}} \tag{96}$$

Since $\gamma < 0$ and quantities fall, the mechanical scale effect is positive, pushing observed pass-through upwards and masking the extent of the exporter's price cut ρ.

A.8 Micro-foundation of the Instrumental Variable Strategy

To motivate the instrumental variable strategy used in Section 4.1, we consider a simple model of a buyer sourcing a variety from multiple suppliers.

Consider a buyer j who demands a composite variety v (e.g., "smartphone screens") for use in production or resale. The buyer aggregates inputs from a set of potential suppliers $\Omega_{j,v}$ using a Constant Elasticity of Substitution (CES) aggregator:

$$Q_{j,v,t} = \left(\sum_{i \in \Omega_{j,v}} \alpha_{i,j,v}^{\frac{1}{\sigma}} q_{i,j,v,t}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \tag{97}$$

where $q_{i,j,v,t}$ is the quantity purchased from supplier i at time t, $\alpha_{i,j,v}$ is a time-invariant quality or preference shifter, and σ is the elasticity of substitution between suppliers.

The buyer chooses quantities $q_{i,j,v,t}$ to minimize the cost of achieving a target aggregate level $Q_{j,v,t}$. The resulting demand for a specific supplier's good is:

$$q_{i,j,v,t} = \alpha_{i,j,v} \left(\frac{p_{i,j,v,t}}{P_{j,v,t}} \right)^{-\sigma} Q_{j,v,t} \tag{98}$$

where $P_{j,v,t} = \left(\sum_{k \in \Omega_{j,v}} \alpha_{k,j,v} p_{k,j,v,t}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ is the ideal price index for the variety.

We assume the buyer's total demand for the variety $Q_{j,v,t}$ is determined by a downstream demand curve with elasticity η:

$$Q_{j,v,t} = D_{j,v,t} P_{j,v,t}^{-\eta} \tag{99}$$

where $D_{j,v,t}$ represents a buyer-specific demand shock (e.g., a seasonal surge in sales, a new contract, or inventory restocking needs).

Substituting the aggregate demand into the supplier-specific demand (98) and taking logs yields:

$$\log q_{i,j,v,t} = \underbrace{\log \alpha_{i,j,v}}_{\text{Fixed Effect}} - \sigma \log p_{i,j,v,t} + (\sigma - \eta) \log P_{j,v,t} + \log D_{j,v,t} \tag{100}$$

Now consider the supply side. The price offered by supplier i is determined by the scale elasticity γ and an idiosyncratic supply shock $\xi_{i,j,v,t}$ (e.g., a machine breakdown or

a batch-specific quality issue):

$$\log p_{i,j,v,t} = \gamma \log q_{i,j,v,t} + \mu_{i,j,v} + \xi_{i,j,v,t} \tag{101}$$

The identification challenge is that $\log q_{i,j,v,t}$ is endogenous to $\xi_{i,j,v,t}$ because a supply shock affects the price, which in turn affects the quantity demanded.

First we directly include controls when you construct an instrument using quantites purchases from the all supplier in other transactions during the same month.

$$IV_{i,j,v,t} = \log \sum_{t \neq t} q_{k,j,v,t'} \tag{102}$$

However, we must directly control for the focal supplier's supply shock $\xi_{i,j,v,t}$ to ensure the instrument is valid. This is achieved by including seller-product-month fixed effects in the regression, which absorb any unobserved shocks at the seller-product-month level (e.g., a supplier-wide productivity shock or a change in input costs for a specific product) that could otherwise correlate with both the instrument and the error term.

Alternatively, the instrument is the sum of quantities purchased from *other* suppliers $k \neq i$:

$$IV_{i,j,v,t} = \log \sum_{k \neq i} q_{k,j,v,t} \tag{103}$$

Relevance: Complements and Substitutes The validity of the instrument relies on the correlation between $q_{i,j,v,t}$ and purchases from other suppliers $q_{k,j,v,t}$. From Equation (100), the quantity purchased from any other supplier k also depends on the common demand shock $D_{j,v,t}$:

$$\log q_{k,j,v,t} = \log \alpha_{k,j,v} - \sigma \log p_{k,j,v,t} + (\sigma - \eta) \log P_{j,v,t} + \log D_{j,v,t} \tag{104}$$

Crucially, the term $\log D_{j,v,t}$ enters positively for *all* suppliers. When the buyer experiences a positive demand shock ($D_{j,v,t} \uparrow$), they increase purchases from supplier i and supplier k simultaneously. This creates a positive correlation between $q_{i,j,v,t}$ and the instrument $IV_{i,j,v,t}$, satisfying the relevance condition.

This logic holds regardless of whether the goods are substitutes ($\sigma > 1$) or complements ($\sigma < 1$).

- **Substitutes ($\sigma > 1$):** If suppliers are substitutes, a price decrease from supplier k would normally reduce demand for supplier i. However, our instrument relies on the *common* shifter $D_{j,v,t}$, which moves both quantities in the same direction,

dominating the substitution effect in the first stage.

- **Complements ($\sigma < 1$):** If suppliers are complements, a price decrease from supplier k increases demand for supplier i. In this case, both the cross-price effect and the common demand shock $D_{j,v,t}$ work in the same direction, reinforcing the positive correlation.

Thus, the instrument is valid for any σ, provided the variation is primarily driven by buyer-level demand shocks rather than idiosyncratic supply shocks to competitors.

Relevance: The Dominance of Common Demand Shocks The validity of using "purchases from other suppliers" (q_{-i}) as an instrument for q_i relies on the existence of a common buyer-level demand fever (D_{it}) that moves purchases from all suppliers in the same direction.

One concern is that if suppliers are substitutes, idiosyncratic price fluctuations could induce negative correlation (buying more from k and less from i). However, in the context of intermediate input sourcing, the "bullwhip effect" implies that high-frequency variation is dominated by inventory "batching" or production runs (Alessandria et al., 2010). When a buyer initiates a production run, they demand inputs from *all* suppliers simultaneously.

Empirically, we observe a strong positive first-stage F-statistic, confirming that the "common demand shock" (complementarity in timing) dominates any idiosyncratic substitution effects. Thus, q_{-i} isolates exogenous variation in q_i driven by buyer scale, allowing us to trace out the seller's supply curve.

Exclusion The validity of the instrument requires that the common demand shock $D_{j,v,t}$ is uncorrelated with the idiosyncratic supply shock of the focal supplier $\xi_{i,j,v,t}$. This holds under the standard assumption that buyer-side demand fluctuations are orthogonal to supplier-specific cost or quality shocks. By controlling for buyer-seller fixed effects (absorbing α and μ) and using the variation from $D_{j,v,t}$ captured by the other suppliers' quantities, we isolate exogenous variation in $q_{i,j,v,t}$ that is uncorrelated with $\xi_{i,j,v,t}$. Furthermore, by including seller-product-month fixed effects, we control for any unobserved shocks at the seller-product-month level (e.g., a supplier-wide productivity shock or a change in input costs for a specific product) that could otherwise correlate with both the instrument and the error term. This ensures that the instrument's validity is not compromised by common shocks affecting all of a supplier's buyers for a given product.

A.9 CES Price Index Calculation

This appendix details the methodology used to construct the aggregate price index. We assume a representative downstream buyer with Constant Elasticity of Substitution (CES) preferences. This framework allows us to account for two distinct sources of price level changes: the changing prices of continuing varieties (the intensive margin) and the entry and exit of new varieties (the extensive margin).

A.9.1 The Theoretical Framework

The unit cost function corresponding to a CES aggregator with elasticity of substitution $\sigma > 1$ is defined as:

$$P_t = \left(\sum_{i \in \Omega_t} b_{it} p_{it}^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{105}$$

where Ω_t is the set of varieties available in period t, p_{it} is the price of variety i, and b_{it} is a taste or quality parameter.

Following Feenstra (1994), the ratio of unit costs between two periods $t-1$ and t can be decomposed into a component tracking goods available in both periods (the common set $I = \Omega_t \cap \Omega_{t-1}$) and a component adjusting for the changing set of varieties.

A.9.2 The Intensive Margin (Sato-Vartia Index)

For the set of goods I common to both periods, the change in the price level is measured using the Sato-Vartia index. This index is "exact" for CES preferences, meaning it contains no substitution bias. It is defined as the geometric mean of price relatives weighted by the logarithmic mean of expenditure shares.

Let s_{it} denote the expenditure share of good i in period t. The weight w_i^* is constructed as follows:

$$w_i^* = \frac{\theta_i}{\sum_{j \in I} \theta_j}, \quad \text{where} \quad \theta_i = \frac{s_{it} - s_{i,t-1}}{\ln s_{it} - \ln s_{i,t-1}} \tag{106}$$

If $s_{it} = s_{i,t-1}$, the logarithmic mean simplifies to $\theta_i = s_{it}$. The price index for common goods, P_{SV}, is then:

$$\ln \left(\frac{P_t^{SV}}{P_{t-1}^{SV}} \right) = \sum_{i \in I} w_i^* \ln \left(\frac{p_{it}}{p_{i,t-1}} \right) \tag{107}$$

A.9.3 The Extensive Margin (Feenstra Adjustment)

To account for entry and exit, we define the λ-ratio, which measures the share of expenditure allocated to the common set of goods relative to the total expenditure in each period. Let I be the set of common goods. We define:

$$\lambda_t = \frac{\sum_{i \in I} p_{it} q_{it}}{\sum_{j \in \Omega_t} p_{jt} q_{jt}} = \sum_{i \in I} s_{it}, \quad \text{and} \quad \lambda_{t-1} = \frac{\sum_{i \in I} p_{i,t-1} q_{i,t-1}}{\sum_{j \in \Omega_{t-1}} p_{j,t-1} q_{j,t-1}} = \sum_{i \in I} s_{i,t-1} \tag{108}$$

Intuitively, $\lambda_t < 1$ implies that new varieties have entered the market in period t, diluting the share of the common goods. Conversely, $\lambda_{t-1} < 1$ implies that varieties present in $t-1$ have exited before period t.

The adjustment factor for product variety is given by:

$$\Lambda_t = \left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma - 1}} \tag{109}$$

Provided that varieties are substitutes ($\sigma > 1$), the entry of new goods ($\lambda_t < 1$) lowers the aggregate price index, reflecting a welfare gain.

A.9.4 The Unified Exact Price Index

Combining the intensive and extensive margins, the total change in the cost of living index (COLI) is:

$$\frac{P_t}{P_{t-1}} = \underbrace{\left(\prod_{i \in I} \left(\frac{p_{it}}{p_{i,t-1}} \right)^{w_i^*} \right)}_{\text{Sato-Vartia Index}} \times \underbrace{\left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma - 1}}}_{\text{Variety Adjustment}} \tag{110}$$

This formula measures of the change in the unit cost function for CES preferences, with both substitution effects and changes in the available product sets.

B Additional Empirical Results and Robustness Checks

B.1 Comparison of Scale Across Methods

Table 17 compares the variety-level scale elasticity estimates (γ) from Section 4.2. Each observation is an HTS 6-digit variety. The table reports regressions of the IV estimates (γ^{IV}), structural estimates ($\gamma^{Structural}$), and shipping cost-based estimates ($\gamma^{Shipping}$) on the baseline OLS estimates (γ^{OLS}). The structural estimates are computed using a method

similar to Feenstra (1994), which assumes locally iso-elastic demand. Columns (1)-(4) report unweighted regressions, while columns (5)-(8) report regressions weighted by the number of observations in each variety. The results show a strong positive correlation between the different estimation methods, particularly between the OLS and IV estimates.

We take these results as evidence that the scale elasticity estimates are robust across different methods and that the OLS and IV estimates are reliable indicators of the true scale elasticity. The strong correlation between these estimates suggests that the differences in the methods used to estimate γ do not significantly alter the overall findings.

In particular, the cost-based estimates are particularly interesting as they provide a different perspective on the scale elasticity, which is not directly influenced by the demand structure or the availability of data on downstream buyer behavior. The strong correlation between the shipping cost-based estimates and the other methods further supports the robustness of the scale elasticity estimates.

Table 17: Comparison of Scale Elasticity Estimates (γ) Across Varieties

	(1) γ^{IV}	(2) $\gamma^{Structural}$	(3) $\gamma^{Structural}$	(4) $\gamma^{Shipping}$	(5) γ^{IV}	(6) $\gamma^{Structural}$	(7) $\gamma^{Structural}$	(8) $\gamma^{Shipping}$
γ^{OLS}	0.834		0.782	0.652	0.906		1.167	1.157
	(0.00974)		(0.0251)	(0.0214)	(0.00764)		(0.0249)	(0.0141)
γ^{IV}		0.683				0.975		
		(0.0244)				(0.0256)		
R^2	0.631	0.155	0.185	0.173	0.766	0.254	0.339	0.602
Weighting	Unweighted				Weighted - Observations			

Notes: This table compares the variety-level scale elasticity estimates (γ) from Section 4.2. Each observation is an HTS 6-digit variety. The table reports regressions of the IV estimates (γ^{IV}), structural estimates ($\gamma^{Structural}$), and shipping cost-based estimates ($\gamma^{Shipping}$) on the baseline OLS estimates (γ^{OLS}). The structural estimates are computed using a method similar to Feenstra (1994), which assumes locally iso-elastic demand. Columns (1)-(4) report unweighted regressions, while columns (5)-(8) report regressions weighted by the number of observations in each variety.

B.2 Recovered Scale Elasticity Correlates

In Table 18, we perform a falsification test of the price discrimination hypothesis. We relax the assumption of a common scale elasticity and instead recover a unique scale elasticity $\gamma_{i,j,v}$ for every buyer-seller pair using the strategy from Section 4.2. If quantity discounts were driven by second-degree price discrimination, we would expect the steepness of the discount schedule to vary with the relative market power of the trading partners.

We regress these relationship-specific elasticities on proxies for market power, including related party status, bilateral volume, and market shares. Standard errors are clus-

tered at the variety level. As shown in Table 18, we find precise null results: none of the market power proxies are statistically significant predictors of $\gamma_{i,j,v}$. This lack of correlation is strong evidence against the price discrimination mechanism. It supports our main interpretation that γ reflects a technological parameter (e.g., logistical cost structure) that is common across transactions, rather than a strategic variable adjusted for specific relationships.

Table 18: Correlates of Buyer-Seller Scale Elasticities

	(1) $\gamma_{i,j,v}$	(2) $\gamma_{i,j,v}$	(3) $\gamma_{i,j,v}$	(4) $\gamma_{i,j,v}$	(5) $\gamma_{i,j,v}$	(6) $\gamma_{i,j,v}$
Related Party	-0.756					
	(0.717)					
Bilateral Quantity: $\log(q_{ij})$		0.00856				
		(0.0964)				
Output Share			-0.159			
			(0.773)			
Input Share				0.0905		
				(0.505)		
Quantity Sold: $\log(q_i)$					-0.00435	
					(0.0144)	
Quantity Bought: $\log(q_j)$						0.00993
						(0.00852)
R^2	0.25	0.25	0.25	0.25	0.0187	0.0187
Within R^2	0.0000101	8.45E-09	5.09E-08	3.34E-08	2.77E-08	0.00000018
Fixed Effects	Buyer-Variety, Seller-Variety				Variety	

Notes: This table uses the OLS strategy from Section 4.2 to recover a scale-elasticity for every buyer-seller pair in the data. Each observation is a buyer-seller-variety scale elasticity. Standard errors are clustered at the variety level.

In 19, we explore the correlates of year-on-year changes in variety scale elasticities from 2016 to 2017. Each observation is a variety scale elasticity. We regress the change in scale elasticities on changes in various factors such as changes in HHI for both buyers and sellers, changes in the number of buyers, sellers, pairs, transactions, and value. Standard errors are clustered at the variety level. We do not find that the changes in scale elasticity are significantly correlated with any of these factors.

B.3 Structural Estimation of Supply and Demand Elasticities

One downside of this reduced form approach is that it doesn't account for the shape of demand (such as the elasticity of substitution), but rather instruments for variation in de-

Table 19: Correlates of Changes in Buyer-Seller Scale Elasticities

	(1) $\Delta\gamma^{OLS}$	(2) $\Delta\gamma^{OLS}$	(3) $\Delta\gamma^{OLS}$	(4) $\Delta\gamma^{OLS}$	(5) $\Delta\gamma^{OLS}$	(6) $\Delta\gamma^{OLS}$	(7) $\Delta\gamma^{OLS}$
Δ HHI Seller	0.639						
	(0.520)						
Δ HHI Buyer		0.359					
		(0.203)					
$\Delta \log(\text{Buyers})$			0.31				
			(0.275)				
$\Delta \log(\text{Sellers})$				0.162			
				(0.194)			
$\Delta \log(\text{Pairs})$					0.149		
					(0.149)		
$\Delta \log(\text{Transactions})$						0.00444	
						(0.0316)	
$\Delta \log(\text{Value})$							-0.00412
							(0.0101)
R^2		0.003	0.001	0.003	0.001	0.001	0.000

Notes: This table explores the correlates of year-on-year changes in buyer-seller scale elasticities (γ) from 2016 to 2017. Each observation is a buyer-seller-variety scale elasticity. We regress the change in scale elasticities on changes in various factors. Standard errors are clustered at the variety level.

mand while simultaneously controlling for supply shocks. An alternative approach is to structurally estimate supply and demand within a relationship over time. Following the trade literature, we can estimate both the within-relationship (inverse) supply elasticity alongside the demand elasticity using panel, which we collect at the monthly level for relationships. This is a simplification from our purely transaction-level approach, but as the estimation require a continuous panel, we compromise at this level of aggregation. Similarly, we also run our analysis at the 6-digit HS code level.

As an alternative, we can put structure on demand (say CES), we can jointly estimate supply and demand. (Following Feenstra (1994), and in an application similar to Hottman, Redding, and Weinstein (2016) which allow for variation in the level of markups), as opposed to the solely supply side exercise with an instrument for demand. There is one downside to this, all demand and supply functions must have fully defined functional forms. We will assume constant elasticity of demand (CES) and a constant elasticity of supply (globally as opposed to locally).

As for identification, fundamentally, there is a timing assumption, shocks in demand between time periods and between trading pairs are orthogonal to shocks on the supply side. This assumption does need a parametrization of both supply and demand to recover residuals.

Traditionally, this analysis is run at the origin-country, destination-country, and product category level (extensions include Broda and Weinstein (2010) and Soderbery (2015)).

Formally, assume the transaction-level demand is CES. Taking logs, taking the time difference and differencing relative to another buyer-seller pair r in the same sector s gives

$$\Delta^{r,t} \ln(q_{vt}) = (-\sigma^s)\Delta^{r,t} \ln(p_{vt}) + \nu_{vt}, \tag{111}$$

where $\Delta^{r,t}$ refers to the double difference.

Next, we assume the following transaction-level pricing equation holds in double-differences form:

$$\Delta^{r,t} \ln p_{vt} = \omega^s \Delta^{r,t} \ln(q_{vt}) + \kappa_{vt}, \tag{112}$$

We assume that the following orthogonality condition holds for each buyer-seller pair:

$$G(\beta_s) = \mathbb{E}_{\mathbb{T}} \left[\nu_{vt}\kappa_{vt}(\beta_s)\right] = 0 \tag{113}$$

where $\mathbb{E}_{\mathbb{T}}$ is the time series expectation and $\beta_s = \begin{pmatrix} \sigma^s \\ \omega^s \end{pmatrix}$.

Re-writing this orthogonality assumption gives:

$$\mathbb{E}_{\mathbb{T}}\left[(\Delta^{r,t}\ln p_{vt})^2\right] = \mathbb{E}_{\mathbb{T}}\left[(\frac{\omega^s}{\sigma^s})(\Delta^{r,t}\ln q_{vt})^2 + (\frac{\omega^s\sigma^s-1}{\sigma^s})\Delta^{r,t}\ln q_{vt}\Delta^{r,t}\ln p_{vt} + \frac{1}{\sigma^s}\nu_{vt}\kappa_{vt}\right], \tag{114}$$

which simplifies to:

$$\mathbb{E}_{\mathbb{T}}\left[(\Delta^{r,t}\ln p_{vt})^2\right] = (\frac{\omega^s}{\sigma^s})\mathbb{E}_{\mathbb{T}}\left[(\Delta^{r,t}\ln q_{vt})^2\right] + (\frac{\omega^s\sigma^s-1}{\sigma^s})\mathbb{E}_{\mathbb{T}}\left[\Delta^{r,t}\ln q_{vt}\Delta^{r,t}\ln p_{vt}\right]. \tag{115}$$

Implementation Implementation of the simultaneous identification of supply and demand is relatively demanding. It requires a continuous time series for multiple buyers and sellers over time.[40] Furthermore the identification of such demand and supply curves are not directly applicable to the idiosyncratic transaction level data. To adapt the data, we make for two adaptations. First, instead of considering products at the root 10-digit level. We aggregate to the 4-digit level. Second, instead of using the raw transactions at the daily level, we aggregate to the monthly level. [41]. Overall, we are left of a time series of buyers and sellers from 2008 to 2016, where the relationship is active for 12 continuous time periods (months in this application). We winsorize price, sales, and quantity changes at the 1% and 99% levels.

There are also a few choices that need to be made for GMM estimation. First is the initial search space. We search over values σ and ω in the space starting by -2 to 2 for ω and 1 to 12 for σ. Second is in the weighting of observations. We largely follow Broda and Weinstein (2006) and weight observations by a function of the count of observations within a trading relationship.

Results Operationally, this technique is more demanding than that in Section 4.2, furthermore it requires a grid search to account for multiple solutions. Following the literature, we do primarily restrict the solution to those with downward sloping demand; but are flexible in terms of supply.

As noted in the literature, there are multiple possible solutions, due to the fact that the solution lies at the intersection of two hyperbolas. We find this common. Typically we find there are two solutions, often one of the solutions is small and with ω near zero and the other solution is broadly falls in the line of our estimates, were $\omega \approx \gamma_v \approx -0.3$.

[40] This technique is infeasible with the domestic data, as data is collected only in a single time period for any origin.

[41] A forthcoming appendix shows that such an adaptation is not problematic under relatively mild conditions, and may realize in an slight underestimate of the supply elasticity

Table 20: Combined Supply and Demand Estimation

	Supply Elasticity (ω)			**Demand Elasticity (σ)**		Correlation
Spec	Median	Mean	Variance	$\frac{1}{mean(\sigma^{-1})}$	$e^{mean(\ln \sigma)}$	(ω,σ^{-1})
1	-0.233	-0.365	0.205	3.425	7.794	0.587
2	-0.618	-0.496	0.183	6.494	11.510	0.465
3	-0.222	-0.357	0.208	3.497	9.976	0.532

Notes: See text for details.

Table 20 shows our results. Fundamentally, ω is similar to the γ estimated, but with a different strategy. Here instead of relying on real-world instrumentation and fixed effects, identification instead relies on the orthogonality of the variance of supply and demand. We consider this.

We first choose the solution with $\sigma > 0$ and with the smallest inverse supply elasticity. In this main specification (Specification 1), finds that a median elasticity of -.233 and a median elasticity of -.365, with a variance of 0.205 across HTS codes. These estimates are broadly in line with our standard identification strategy and illustrates the robustness of both the instrumentation strategy and well as the consistent identification of demand.

In terms of the demand side, there is a wide right tail in σ, with some point estimates trailing off towards infinity. Some researcher use ad-hoc trimming of the underlying data to get that under control. Instead we report the inverse of the mean of the inverse. This transformation (and the related logramthic and exponential) transformation, finds a mean demand elasticity of approximately of $\sigma = 3.4$, which is broadly in line with the literature.[42] Furthermore, the supply elasticity ω has a correlation of 0.47 with that we recover using the reduced form methodology (the parameter γ).

Overall, this is broadly consistent with the supply elasticities estimated in the more reduced-form exercise Section 4.2, as well as firm-demand estimates of σ from Hottman et al. (2016).[43]

B.4 Alternative Time Periods

In table 21, we replicate 11 for the year 2019. We use the same specifications as in Table 11, but we only include transactions from 2019. The results show that the pass-through of tariffs to observed prices is still significant, but the magnitude of the effect is lower

[42]This is similar to the median, but which we have not disclosed from the Census.

[43]It also lines up broadly with Broda and Weinstein (2006) and others at the level of substitution between countries

compared to the 2017-2018 period. However the gap between the effect for raw unit values (column 1) and the effect for quantity-adjusted values (columns 2 and 3) is similar compared to the 2018 tariffs.

Table 21: Transaction-level Tariff Pass-Through: Within Relationships: 2019 Only

	(1) $\log(p)$	(2) $\log(\tilde{p}^{OLS})$	(3) $\log(\tilde{p}^{IV})$	(4) $\log(pq)$	(5) $\log(q)$	(6) $\log(p)$
log(1 + Tariffs Applied)	0.865	0.667	0.648	-0.0563	-0.921	
	(0.0284)	(0.0212)	(0.0366)	(0.0324)	(0.0432)	
log(1+ Tariffs Statutory)						0.975
						(0.340)
R^2	0.953	0.956	0.991	0.715	0.897	0.954
Within R^2	0.00104	0.00105	0.000522	0.00000271	0.00049	0.00000124
Fixed Effects	Buyer-Seller-Variety, Variety-Year-Month					

Notes: This table reports transaction-level tariff pass-through estimates for continuing buyer-seller-variety relationships in 2019. All specifications include buyer-seller-variety and variety-year-month fixed effects. Standard errors are clustered at the relationship level. Column (1) shows the pass-through of applied tariffs (duty paid) to observed transaction prices. Columns (2) and (3) report pass-through to quantity-adjusted prices ($\tilde{p}$), constructed using the OLS and IV estimates of the scale elasticity γ from Section 4.2, respectively. These columns isolate the change in the price schedule from compositional effects. Column (4) reports the effect on transaction value, while Column (5) shows the response of transaction quantity. Column (6) shows the pass-through of statutory tariffs.

Table 22 reports the results of estimating the aggregate pass-through of tariffs to observed unit values, quantity-adjusted unit values, and other relevant variables for the period 2017-2018. We drop the 2019 tariffs, which did not have long to be in before the COVID-19 shock. The table includes fixed effects at the product-origin-month level and clusters standard errors at the product-origin level. The coefficients in column (1) represent the reduced-form pass-through of tariffs to observed unit values, while column (2) reports pass-through to quantity-adjusted unit values, constructed by netting out transaction-level scale effects using the estimated γ from Section 4.2. Column (3) decomposes the scale-adjusted price into the composite of markups and marginal costs $\mu\tilde{c}$, isolating the change in seller pricing behavior net of the mechanical tariff effect. Column (4) reports the residual quantity composition effect $\tilde{q}$. Column (5) shows the effect on total transaction values pq, column (6) reports the quantity response, column (7) shows the change in average transaction size q/T, and column (8) reports the change in the number of transactions T. Column (9) presents the first-order approximation of scale-adjusted pass-through computed as $-\gamma\times$ (coefficient from column 7), using $\gamma = 0.29$ from Section 4.2.

Table 22: Monthly Aggregate Pass-Through, 2017-2018

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	$\log p$	$\log \tilde{p}$	$\log \mu\tilde{c}$	$\log \tilde{q}$	$\log pq$	$\log q$	$\log(q/T)$	$\log T$	$\widehat{\log \tilde{p}}$
log(1+Tariffs)	1.01	0.593	-0.417	0.427	-1.493	-2.502	-1.881	-0.621	0.540
	(0.0376)	(0.0149)	(0.0149)	(0.0280)	(0.0395)	(0.0548)	(0.0472)	(0.0229)	
r^2	0.908	0.878	0.878	0.902	0.87	0.9	0.892	0.904	
r^2 within	0.0003	0.0006	0.0006	0.0001	0.0006	0.0009	0.0007	0.0003	
Fixed Effects	Variety-Year-Month, Variety-Country Origin								

Notes: This table reports aggregate tariff pass-through estimates at the product-origin-month level over 2017-2018. Column (1) shows the reduced-form pass-through of tariffs to observed unit values. Column (2) reports pass-through to quantity-adjusted unit values $\tilde{p}$, constructed by netting out transaction-level scale effects using the estimated γ from Section 4.2. Column (3) decomposes the scale-adjusted price into the composite of markups and marginal costs $\mu\tilde{c}$, isolating the change in seller pricing behavior net of the mechanical tariff effect. Column (4) reports the residual quantity composition effect $\tilde{q}$. Column (5) shows the effect on total transaction values pq. Column (6) reports the quantity response. Column (7) shows the change in average transaction size q/T. Column (8) reports the change in the number of transactions T. Column (9) presents the first-order approximation of scale-adjusted pass-through computed as $-\gamma\times$ (coefficient from column 7), using $\gamma = 0.29$ from Section 4.2. All specifications include product-time and product-origin fixed effects. Standard errors are clustered at the product-origin level.
XXX - check column (2) SE

B.5 Decomposition of Aggregate Price Variation

We can then run regressions to see if on this, along the vein of the decomposition in Section 3 to understand what correlates to aggregate relationship-level prices, as those studied by Alviarez et al. (2023) and Kamal and Sundaram (2016):

$$\log \tilde{p}_{i\to j,v} = \theta \log q_{i\to j,v} + FE_{i,v} + FE_{j,v} + \epsilon_{ij}. \tag{116}$$

Table 23: Relationship-Level Cross Sectional Price Differences

	(1)	(2)	(3)	(4)
	$\log p$	$\log \tilde{p}$	$\log p$	$\log \tilde{p}$
$\log q$	-0.221	-0.074	-0.136	-0.020
	(0.002)	(0.001)	(0.003)	(0.001)
r^2	0.391	0.236	0.906	0.89
Within r^2	0.223	0.052	0.135	0.006
Fixed Effects	Country-HTS10		Buyers-HTS10	
			Seller-HTS10	

Notes: See text for details.

Table 23 highlights these regressions in two forms. The first using raw price data and the second using our residual price $\tilde{p}$, which controls for the transaction-level quantity discount. Column (1) shows that with product and source fixed effects, the aggregate quantity between a seller and buyer, accounts for 23% of the price variation, with an aggregate elasticity θ = -0.22. Column (2) shows that accounting for transaction level quantities implies that aggregate quantities only explain 5% of the price variation. However this accounting implicitly makes assumptions about the unique nature of a product.

Columns (3) and (4) include the most demanding set of fixed effects, buyers-variety and seller-variety. Column (3) shows that there is a relatively significant correlation between raw prices and aggregate quantities, with θ = -.137, explaining 13.5% of the variation. However, when controlling for transaction level quantities in (4), there is an extremely small effect, with almost no relationship between the quantity bought and the price paid.[44]

However, these regressions are purely a decomposition, they aren't reflecting a supply curve, but rather the equilibrium output of millions of relationships in the cross-section. To do so, we turn to a first difference strategy, that exploits changes in underlying demand to estimate a relationship level supply curve, both with and without accounting for economies of scale in individual transactions.

[44]We additionally run regressions on importer and exporter shares and volumes, as in section 3, but find they have minimal additional explanatory power as before, once relationship-level quantities are controlled for.

C Domestic Trade Data

C.1 CFS Data Descriptions

For domestic trade, we consider the U.S. Census and Department of Transportation Commodity Flow Survey (CFS), focusing on 2012. This survey asks US-based manufacturing and wholesale establishments about their transactions, recovering the weight, value, method of transport, and the final shipping destination. Products are classified by a 5-digit Standard Classification of Trade Goods (SCTG) code. There are three limitations to this data, compared to the import data. First, we do not observe the identity of the receiving firm, but rather only their zip-code or state, which we code as the final buyer (thus a buyer is a product-location combination). Second, for goods that are not sold by weight, we do not see a measure of quantity, as such, we drop SCTG codes for goods not sold by weight.

This survey of approximately 100,000 establishments (out of 700,000 potential establishments) covers locations that originate 80% of domestic manufactured good shipping. Only a subset of transactions are recorded to keep reporting requirements tractable, resulting in about 5 million transactions in 2012. We decompose prices for 2012, but carry out further analysis every 5 years from 1997 to 2017.

Table 24: Variance Decomposition of Domestic Trade Data

	Specification		**Variance Decomposition**			
	Controls	Fixed Effects	Controls	Fixed Effects	Covariance	Residual
1	$\log q$		32.5%			67.5%
2	$\log q$ x Variety		37.6%			62.4%
3		Seller-Buyer-Variety		73.7%		26.4%
4	$\log q$	Seller-Buyer-Variety	18.6%	47.3%	12.0%	22.1%
5	$\log q$ x Variety	Seller-Buyer-Variety	20.4%	45.6%	12.4%	21.6%

Notes: We decompose transaction-level price variation after demeaning log(price) by product code (SCTG 5-digit). Only products where quantities are regularly denominated in weight are included. Sellers are designated at the establishment and product code level. Buyers are designated at the destination state and product code level. For consistency, the sample is fixed to remain constant over the sample period. See the text for full details and specification.

C.1.1 Decomposing Domestic Trade Prices

We conduct a similar exercise with the US domestic Commodity Flow Survey (CFS). The data are limited in certain aspects, particularly because they represent a highly restricted

sample from the buyer's perspective.[45] We broadly find similar trends to the international trade import data.[46]

Data on quantities is restricted to material sold by weight, as non-weight quantities are not recorded (as in international trade). We subset analysis only to goods sold traditionally by weight.[47] Unlike in US import data, data is not collected on the identity of the buying firm, only the identity of the seller is known. We only know the address of the firm receiving shipment. We use the ZIP code of the recipient firm as a proxy for the buyer. Product codes are substantially more aggregated. A 5-digit SCTG code is largely comparable to a 4-digit HTS import code or 4-digit HS trade code.

In Table 24, Row (1) replicates the exercise to decompose the variance of prices on the size of a transaction (all demeaned and residualized at the SCTG 5-digit level). Transactions are identified with a seller, product code, shipment date and buying zip code tuple.

Specification (1) finds that 32.5% of all price variation is due to a log-linear quantity relationship. Allowing this relationship to vary across SCTG 5-digit products increases this to 37.6% of all variation. The share of prices explained by quantity is indeed a bit lower than those in the international trade data, however that data has better measures of product categories and direct measures of quantities. Regardless of the data generating process, within tightly defined product groups, a simple log(q) explains between 30-40% of all price variation.

Specifications (3)-(5) attempt to see if this relationship is purely reflective of the buyer-seller pair, or if there is any explanatory power due to transaction-level quantities. Broadly, as with international trade, the results are similar.[48]

C.2 Supply Estimation: Domestic Data

We repeat a version of the supply isolation using the domestic shipment data from the CFS. The available data makes it hard to directly control for supply shocks, but an aspect of the data collection process works in our favor. In particular, the government questionnaire recovers a sample of shipments, not throughout the year, but within a tightly defined window (typically a week or so). This mechanically controls for supply-time fixed effects. As such, all supply shocks (after controlling for the relationship), must differentially affect some, but not all other customers. Our primary identification threat is

[45] This is important since many types of production function estimation assume away price variation in input data.

[46] This analysis is replicable in public use data that aggregates regions and SCTG codes.

[47] Loosening this restriction does not significantly change our results.

[48] We replicate this analysis using the 2012 public use data in a forthcoming appendix.

unobserved qualities across shipments.

As before, we create an instrumental variable that is correlated with a demand shock, but uncorrelated with the current supplier (for which we have implicitly a supplier-time control). Following Equation 10, we create three instruments. The first looks at how much of a product is bought from a firm by a destination state, excluding the focal order. This naturally is at roughly the month level, as orders from firms are clustered due to sampling. The second considers how much is bought by a destination state within a month-product, excluding the focal transaction. The last, does the same, but excludes all transactions from the original focal firm.

Table 25 produces a version of Table 2-3 for our domestic data in 2012. Column (1) illustrates in the simply demeaned data a correlation, where γ is equivalent to -30.4%; broadly in line with the import data. However, looking within a buyer-seller-product relationship in column (2), the same γ is -0.23. Columns (3) and (4) repeat this exercise, but allow γ_v to vary across each of the 5-digit SCTG products. Mean estimates of γ_v are -0.295 and -.206, with relatively tight variances across different products.

Turning to our instrumentation strategy to control for unobserved product quality, as well as the limited possibility of supplier shocks, columns (5)-(6) display the results. Column (5) replicates column (2), but with the same sample as observations with our first instrument. Column (6)-(8) show that all three instrumentation strategies yield strong predictive first stages and have γ_v between -0.21 and -0.26.

Broadly, within-relationship inverse supply elasticities are around -0.28 in the trade data and -0.24 in the domestic trade data. However, the goods traded are different and product level comparisons may yield even more consistent estimates.

C.2.1 Aggregate Scale Economies Domestic Trade

There do not appear to be aggregate scale economies (or cross-sectional or time series effects of changes in aggregate order volumes and indication of unilateral or bilateral market power) in international imports. What about domestic trade? We conduct the same exercise, first a decomposition, then a regression controlling for supply side effects with an instrument.

There are two downsides to the domestic trade data. First, it is not a time series, severely limiting both controls and potential instruments. However, even in the cross-section, we can control for relationship fixed effects. Furthermore, the domestic data has one advantage over the trade data: for a single seller, we see a wider swath of buyers and can directly control for some bilateral effects.

Table 25: Recovering Quantity Discounts: Domestic Data

	(1)	(2)	(3)	(4)
	log p			
log q	-0.304	-0.23		
	(0.00146)	(0.00174)		
log q x SCTG			-0.295	-0.206
			[0.136]	[0.0246]
r^2	0.325	0.779	0.376	0.784
Within r^2	0.325	0.159	D	D
Fixed Effects				
Relationship		✓		✓

	(5)	(6)	(7)	(8)
	log p			
log q	-0.22	-0.209	-0.256	-0.261
	(0.00197)	(0.00170)	(0.00423)	(0.00508)
r^2	0.788	0.155	0.151	0.15
Within r^2	0.155			
Fixed Effects				
Relationship	✓	✓	✓	✓
First Stage F-Stat		74000	69000	11000
Instruments		IV1	IV2	IV3

Notes: Round parenthesis represent standard deviations. Square brackets represent the variance across SCTG 5-digit codes for estimates. The mean estimate is displayed over it. Demeaning refer to process of regularizing all variables by product code (HTS 10-digit) fixed effects. Sellers are designated at the establishment and product code level. Buyers are designated at the domestic destination state and product code level. See the text for full details and specification. D denotes variables not yet disclosed from census. Standard errors are clustered by the buyer-seller pair.

In Table 26, columns (1) and (2), we show that while a substantial portion of the variation in aggregate prices log p can be correlated with total volumes log q, once transaction level discounts are stripped out, there are no correlations. Columns (3) and (4) add significant buyer-product and seller-product fixed effects and show that this relationship is unchanged.

However, to show causality, we need a shifter of demand, not some potential bilateral shifter of supply. Columns (5) and (6) implement a simplified version of the instrument used in the import data, at the destination, how much overall demand exists, without including the focal source. Essentially a location that buys lots of car parts will have

more demand for car parts, even excluding the focal source location. Further to control for geographic proximity, we control for distance (as measured using road distance as computed by the US Census and Department of Transportation).

Without controlling for transaction level quantity discounts, $\theta = -.11$. However, controlling for transaction level quantity discounts, $\theta = .02$, close to constant returns to scale.

Table 26: Decomposing Relationship Prices - Domestic Trade

	(1) $\log p$	(2) $\log \tilde{p}$	(3) $\log p$	(4) $\log \tilde{p}$	(5) $\log p$	(6) $\log \tilde{p}$
$\log q$	-0.15	0.0188	-0.15	0.0288	-0.112	0.024
	(0.00607)	(0.00634)	(0.00609)	(0.00628)	(0.00605)	(0.00609)
log(Distance)			-0.0329	-0.0161	-0.0154	0.0139
			(0.00310)	(0.00280)	(0.00353)	(0.00384)
r^2	0.823	0.947	0.86	0.957	0.112	0.12
Within r^2	0.123	0.00219	0.12	0.00481		
First Stage F-Stat					396.6	396.6
Fixed Effects	Seller, SCTG5		Buyer-SCTG5, Seller-SCTG5			
Instruments					✓	✓

Notes: This comes from a regression of change in aggregates quantities on a measure of price. Odd columns use aggregate prices. Even columns follow the correction in equation X, and adjust each transaction's price for the quantity sold, using the aggregate relationship from (116). Columns (1) through (6) use fixed effects in increasing order of stringency. The last two columns instrument for the change in quantity using the shift-share instrument detailed in the text. Standard errors are clustered by SCTG 5-digit product.