

TWO-SIDED MARKET POWER IN FIRM-TO-FIRM TRADE^{*}

Alviarez, V.
IADB

Fioretti, M.
Bocconi

Kikkawa, K.
UBC Sauder and NBER

Morlacco, M.
USC

(First draft: December 2022)
May 1, 2026

Abstract

We develop and estimate a structural model of bargaining in firm-to-firm trade to study the determinants of tariff pass-through. The model features oligopoly and oligopsony power and yields analytical expressions for bilateral markups and pass-through based on two sufficient statistics: the supplier's share in the buyer's purchases and the buyer's share in the supplier's output. Using U.S. import data, we find substantial importer bargaining power and steep export supply curves. These primitives imply that cost changes, rather than markup adjustments, dominate pass-through, accounting for the bulk of incomplete pass-through of the 2018 U.S. tariffs and its heterogeneity across buyer-supplier links.

^{*}We thank the editor Arnaud Costinot and three anonymous referees for constructive comments and suggestions. For helpful conversations and discussions, we thank Rodrigo Adao, Pol Antràs, Ariel Burstein, Julia Cajal-Grossi, Thomas Chaney, Keith Head, Ali Hortacsu, Michael Irlacher, Oleg Itshkoki, Amit Khadelwal, Yuriy Gorodnichenko, Guillermo Marshall, Kiminori Matsuyama, Thierry Mayer, Isabelle Mejean, Marc Melitz, Scott Orr, Ina Simonovska, Felix Tintelnot, Trevor Tombe, Gustavo Ventura, and Jing Zhang. We also thank participants at seminars at ASU, UC Berkeley, Boston College, Boston University, CREST, Dartmouth, Duke, FED Board, Georgetown, Harvard, UC Irvine, ITAM, JGU Mainz, Keio, LMU, LSE, MIT, Nottingham, NYU, Osaka, Oxford, Peking, Penn State, Princeton, Queens, UC San Diego, Sciences Po, Seoul National, Tokyo, UBC, UCL, USC, UToronto, Wisconsin Madison, Yokohama National, and Yonsei. Siying Wang provided excellent research assistance. Any views expressed are those of the authors and not those of the U.S. Census Bureau. 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 under FSRDC Project Number 2109. (CBDRB-FY25-P2109-R12520). Alviarez and Kikkawa are supported in part by funding from the Social Sciences and Humanities Research Council. All errors are our own. Email: valviarezr@iadb.org (Alviarez); fioretti.m@unibocconi.it (Fioretti); ken.kikkawa@sauder.ubc.ca (Kikkawa); morlacco@usc.edu (Morlacco).

1 Introduction

International trade is organized around a network of buyers and suppliers that transact repeatedly across borders and jointly generate rents along global supply chains. How these rents are divided, and who bears cost shocks, has first-order implications for the incidence of trade policy. Yet canonical trade models abstract from bilateral relationships, thereby missing key mechanisms of shock transmission. We develop and estimate a structural model of firm-to-firm trade where two-sided market power determines the level and heterogeneity of relationship-level markups and pass-through. Applied to the 2018 U.S. tariffs, our model explains within-relationship pass-through as the outcome of bilateral bargaining between importers and exporters.

Pass-through of a cost shock reflects two distinct forces. The first is a *markup channel*, whereby the negotiated price-cost margin adjusts to the shock, reallocating surplus within the match. The second is a *cost channel*: at the higher effective price, large buyers contract demand, moving suppliers down their supply curves and lowering marginal cost. The relative strength of these forces is ultimately an empirical question that depends on underlying primitives. A central primitive is the inverse export supply elasticity, which not only governs oligopsony power but also shapes the balance between the two channels of transmission through the slope of exporter’s marginal cost curve. With bilateral bargaining, these elasticities are match-specific, generating systematic heterogeneity in pass-through across trading relationships that models with atomistic buyers or suppliers are unable to capture.

We estimate the model using pre 2017 U.S. import data and apply it to analyze the 2018 tariff episode. The estimates indicate strong importer bargaining power and steep export supply curves, consistent with substantial oligopsony power. At these parameter values, the cost channel drives tariff pass-through, with marginal cost adjustments dominating markup responses. Pass-through varies systematically across relationships and declines with buyer share, so larger importers face lower pass-through through a higher residual inverse supply elasticity. In the aggregate, pass-through is incomplete at 67 to 73 percent, implying that exporters absorb roughly one third of the tariff burden, through reductions in scale and quasi-rents. This match-level perspective contrasts with product-level studies that find near complete pass-through of the 2018 tariffs,¹ while reduced-form evidence supports the model predictions relative to alternatives without bargaining or upward-sloping export supply.

The model extends the network-based oligopoly framework of [Dhyne et al. \(2022\)](#) to incorporate two-sided market power and decreasing returns to scale, which generate upward-sloping

¹We reconcile the match level and product level evidence later in the introduction.

export supply and capture lock-in effects in global supply chains (Antràs, 2015). Prices are set through bilateral Nash bargaining, with outside options given by each party’s alternative incumbent relationships. The equilibrium yields analytical expressions for bilateral markups and pass-through as functions of two sufficient statistics: the *supplier share*, measuring the exporter’s concentration among the buyer’s input expenditure, and the *buyer share*, measuring the importer’s concentration among the exporter’s total input supply.

Our framework offers a unified theory of exporter and importer market power in firm-to-firm trade. On the exporter side, rents arise from product differentiation and supplier concentration, expressed as an *oligopoly markup* over marginal cost that rises with the exporter’s supplier share, as in standard models (e.g. Atkeson and Burstein, 2008). On the importer side, buyer power takes two distinct forms. Positive bargaining power gives a buyer countervailing power: the ability to extract a share of supplier rents, without affecting their scale. A positive inverse export supply elasticity further generates quasi-rents, resulting in an *oligopsony markdown*: a price-cost ratio below one that declines with the importer’s buyer share, as larger buyers capture a greater share of those quasi-rents through bargaining.² Under two-sided market power, the equilibrium price-cost margin is a convex combination of the oligopoly markup and the oligopsony markdown, with bargaining power as weight.

We study pass-through within existing buyer-supplier relationships: local price responses holding fixed all other match prices and general-equilibrium variables, and abstracting from longer-run adjustments driven by relationship entry and exit. A novel insight is that the two primitives that govern oligopsony power, namely the inverse export supply elasticity and bargaining power, also govern the pass-through determinants: the former determines whether pass-through reflects cost or markup adjustments; the latter shapes the nature of markup adjustment. When the inverse supply elasticity is high, demand contractions produce large movements in exporters’ marginal cost, and the cost channel drives most of the pass-through, regardless of bargaining power; when it is low, the cost channel is more muted and markup adjustments dominate. Within the markup channel, bargaining power governs the nature of strategic interaction: low buyer bargaining power implies strategic complementarities, which dampen pass-through (e.g., Amiti et al., 2014, 2019a); high buyer bargaining power implies strategic substitutabilities arising from oligopsony power, which amplify it.

We take the model to U.S. transaction-level customs data covering the 2018 tariff episode. Our main data source is the U.S. Census Longitudinal Firm Trade Transaction Database

²In standard oligopsony models, oligopsony power operates through quantity distortions: buyers restrict demand to reduce the price they face. In our model, quantities always lie on buyers’ demand curves and there are no such distortions in equilibrium; the inverse supply elasticity nonetheless governs oligopsony power through the bargaining channel. See Section 2 for details.

(LFTTD), which records values and quantities of import transactions at the importer-exporter-product level. Our main sample focuses on repeated, arm’s-length relationships involving the exchange of intermediate and capital goods. Identification exploits cross-sectional price variation across buyers sourcing from the same exporter, product, and year: comparing prices across buyers within an exporter-product-year cell differences out exporter-level cost variation, so that remaining price differences reflect bilateral price-cost margins, which the model maps to observed market shares and unobserved parameters. We estimate importer bargaining weights of approximately 0.8, four times those of foreign suppliers, and a (short-run) returns-to-scale parameter around 0.5, implying an average residual inverse export supply elasticity of 0.3 for the average link.³ Under constant returns to scale, this elasticity would be zero and both oligopsony power and the cost channel would vanish.

Price and pass-through patterns in the data closely align with the model’s predictions. Controlling for marginal costs, prices rise with the exporter’s supplier share and fall with the importer’s buyer share, consistent with two-sided market power. Out of sample, the model closely replicates the 2018 tariff episode, matching both the average level of pass-through (78% in the model, against 87% in the data) and its cross-sectional gradient: a one-unit increase in buyer share lowers pass-through by 0.42 in the model, against 0.41 in the data. Quantity responses point in the right direction but are estimated less precisely.

Specifications imposing constant returns to scale generate a near-zero interaction between pass-through and buyer share and are readily rejected. A specification retaining decreasing returns but setting bargaining power to zero performs comparably on this gradient: consistent with the theory, under strong decreasing returns the cost channel dominates and the gradient is insensitive to bargaining power. The two specifications differ, however, on average pass-through. Strong importer bargaining power implies strategic substitutabilities among buyers that partially offset the cost channel, while zero bargaining power implies strategic complementarities that reinforce it, pushing pass-through further below one. The baseline model outperforms its no bargaining counterpart in matching average pass-through (78% vs. 72%), as it lies closer to the empirical point estimate, although both fall within the confidence interval.

We use the estimated model to study aggregate pass-through in firm-to-firm trade and decompose tariff incidence within these relationships. Weighting by initial import shares, aggregate pass-through is substantially incomplete, with elasticities between 66 and 73%, below the unweighted average of 87% because larger (high-share) buyers experience the lowest pass-

³This estimate is below standard returns-to-scale estimates from the production function literature, reflecting a fundamentally different object: short-run supply constraints within existing relationships rather than long-run technological returns to scale. We discuss this in Section 4.

through. The adjustment operates almost entirely through the cost channel: shutting down markup elasticities leaves aggregate pass-through largely unchanged, whereas eliminating cost elasticities raises it significantly. Roughly one third of the tariff burden is thus borne by foreign exporters, absorbed not through compressed margins but through lost quasi-rents as lower demand moves exporters down their marginal-cost schedules.

A related literature documents near complete pass-through of the 2018 tariffs at the product level, implying that U.S. buyers bore most of the tariff burden (Amiti et al., 2019b, 2020; Flaaen et al., 2020; Fajgelbaum et al., 2020; Cavallo et al., 2020). Our finding of incomplete pass-through within continuing relationships helps reconcile these results. Product-level estimates combine within-relationship price adjustment with changes in the set of active trading partners. Our results indicate that a full account of aggregate incidence requires understanding how the set of trading relationships responds to tariffs.

Related Literature Our paper contributes to three main strands of the literature.

First, the paper relates to the growing literature on importer market power. Early contributions emphasize the role of large buyers in shaping supplier outcomes (Bernard et al., 2019, 2022), while more recent firm-to-firm evidence documents systematic price discrimination across buyers, consistent with buyer power (Fontaine et al., 2020; Huang et al., 2024; Ignatenko, 2024; Macedoni and Mattana, 2024). Despite this evidence, formal theories of buyer power in international trade remain limited.⁴ Exceptions include Morlacco (2019), who estimates substantial oligopsony power among French importers, and Atkin et al. (2024), who show that importer-supplier bargaining shapes the effects of trade policy.

We develop a tractable framework that nests Dhyne et al. (2022) and extends it to incorporate importer oligopsony power via bilateral bargaining and upward-sloping export supply. These ingredients are quantitatively central: together they account for cross-sectional markup comovements and the level and heterogeneity of pass-through.

Second, the paper relates to the literature on determinants of incomplete pass-through using micro-level data. A large literature attributes incomplete pass-through to strategic markup adjustment under oligopoly (e.g., Amiti et al., 2014, 2019a),⁵ while an alternative line emphasizes cost-side adjustment via nontraded distribution and retail costs in the importing country (Goldberg and Hellerstein, 2013). Our cost channel instead arises upstream, from

⁴Alternative explanations include endogenous matching, whereby more productive buyers source from a larger set of suppliers and intensify upstream competition (Huang et al., 2024), and supplier-side price discrimination, under which buyers with stronger outside options obtain lower markups (Ignatenko, 2024).

⁵See also Berman et al. (2012); Auer and Schoenle (2016); Garetto (2016).

exporters’ decreasing returns and demand contractions by large buyers.⁶ Which channel prevails is state-dependent: in our setting, where short-run supply constraints bind, the cost channel dominates, while the standard markup approach is nested as a special case under constant marginal costs and applies naturally to longer-run horizons.

Finally, we contribute to the macro literature on shock propagation in production networks. A large body of work shows how input-output linkages affect the propagation of idiosyncratic shocks (Acemoglu et al., 2012; Grassi, 2018).⁷ A common feature of this literature is that prices are either taken as given or governed by fixed markups, so shocks propagate primarily through quantities. Acemoglu and Tahbaz-Salehi (2025) allow for bilateral bargaining to study the consequences of supply chain disruptions. In their framework, however, price distortions play no role. We complement this approach by characterizing how unit prices respond to cost shocks under bilateral bargaining with decreasing returns, and showing that price adjustment itself is a quantitatively important transmission channel.

The remainder of the paper is organized as follows. Section 2 develops the theoretical framework and derives the results that guide the empirical analysis. Section 3 describes the data and presents reduced form evidence on prices and pass-through. Section 4 outlines the structural estimation strategy and reports the results. Section 5 examines aggregate pass-through in firm-to-firm trade. Section 6 concludes.

2 Theoretical Framework

This section develops a tractable pricing theory of firm-to-firm trade. We focus on how bilateral bargaining shapes prices and quantities within existing trade relationships, taking the network structure as given to isolate short-run tariff transmission. The framework is static and considers single-product negotiations; repeated interactions and multi-product relationships are incorporated in reduced form through the estimated bargaining parameters (see Section 4). We also abstract from nominal rigidities, such as fixed-price contracts or currency denomination, which appear to play a limited role in tariff pass-through.⁸ All proofs and derivations are in Appendix A.

⁶Prior work on pass-through under bargaining (Gopinath and Itskhoki, 2011; Goldberg and Tille, 2013) imposes constant marginal costs and therefore abstracts from oligopsony and the cost channel.

⁷Empirical evidence confirms these mechanisms. See, for example, Barrot and Sauvagnat (2016), Boehm et al. (2019), and Carvalho et al. (2021).

⁸For example, Amiti et al. (2020) shows similar short- and long-run pass-through for the 2018 U.S. tariffs.

2.1 Environment

We focus on the relationship between exporter i and importer j of an intermediate input. We denote by $\mathcal{Z}_i$ the set of importers connected to exporter i , and by $\mathcal{Z}_j$ the set of exporters connected to importer j . These sets vary across firms and are taken as given.

Exporters and Supply. Exporter i supplies a differentiated input to importers in $\mathcal{Z}_i$. Total exported quantity is $q_i = \sum_{j \in \mathcal{Z}_i} q_{ij}$, where q_{ij} denotes importer j 's demand.

The exporter's marginal cost is

$$MC_i(q_i) = k_i q_i^{\frac{1-\theta}{\theta}}, \quad (2.1)$$

where $\theta \in (0, 1]$ governs returns to scale and k_i captures exogenous cost shifters, such as productivity or foreign wages. When $\theta < 1$, the technology exhibits decreasing returns, so marginal cost rises with output. When $\theta = 1$, marginal cost is constant and equal to k_i . The implied total cost function is $TC_i(q_i) = \theta k_i q_i^{\frac{1}{\theta}}$, so average cost equals θ times marginal cost.

Importers and Demand. Importer j produces a final good q_j by combining a foreign input (q_j^f) with a domestic input (q_j^d) according to the production function

$$q_j = \varphi_j (q_j^f)^\gamma (q_j^d)^{\varrho-\gamma}, \quad (2.2)$$

where φ_j denotes productivity, γ the output elasticity of foreign inputs, and ϱ the returns-to-scale parameter.⁹

The foreign input consists of a bundle of input varieties supplied by the exporters in $\mathcal{Z}_j$ and aggregated according to a CES technology:

$$q_j^f = \left(\sum_{i \in \mathcal{Z}_j} \varsigma_{ij} q_{ij}^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}}, \quad (2.3)$$

where ς_{ij} is a demand shifter, and $\rho > 1$ is the substitution elasticity across input varieties.

Downstream, importer j sells its final good in a monopolistically competitive market and faces CES demand with elasticity $\nu > 1$. Revenue is therefore given by

$$R_j(q_j) = D_j q_j^{\frac{\nu-1}{\nu}}, \quad (2.4)$$

⁹Since domestic input purchases are not observed, we impose a unit elasticity of substitution between foreign and domestic inputs, which fixes the cost share of q_j^d and q_j^f . Under this restriction, q_j^d can be interpreted as a constant-returns aggregator of primary factors such as labor or domestic intermediates.

where D_j denote an exogenous demand shifter.

2.2 Bargaining Protocol and Gains From Trade

Each bilateral interaction between importer j and exporter i unfolds in two stages. In the first stage, for a given output target q_j , importer j chooses foreign inputs $\{q_{\ell j}\}_{\ell \in \mathcal{Z}_j}$ and domestic input q_j^d to minimize total cost, taking input prices as given:

$$\min_{\{q_{\ell j}\}, q_j^d} \sum_{\ell \in \mathcal{Z}_j} p_{\ell j} q_{\ell j} + p_j^d q_j^d \quad \text{s.t.} \quad \varphi_j(q_j^f)^\gamma (q_j^d)^{\rho-\gamma} \geq q_j. \quad (2.5)$$

Solving this problem yields the standard CES demand for input varieties:

$$q_{ij} = \varsigma_{ij}^\rho \left(\frac{p_{ij}}{p_j^f} \right)^{-\rho} q_j^f, \quad (2.6)$$

where q_j^f is the foreign input aggregator and p_j^f the associated price index, satisfying $p_j^f q_j^f = \gamma c_j q_j$, with $c_j = MC_j(q_j)$ denoting importer j 's marginal cost.

In the second stage, the two parties bargain over the price p_{ij} , taking the demand schedule (2.6) as given.¹⁰ Operating profits are given by $\pi_{ij}^i = \sum_{k \in \mathcal{Z}_i} p_{ik} q_{ik} - TC_i(q_i)$ for the exporter and $\pi_{ij}^j = R_j(q_j) - \sum_{\ell \in \mathcal{Z}_j} p_{\ell j} q_{\ell j} - p_j^d q_j^d$ for the importer, where $TC_i(q_i)$ and $R_j(q_j)$ are defined in Section 2.1. If bargaining fails and link (i, j) is severed, each party falls back on profits from its remaining matches, denoted $\tilde{\pi}_{(-j)}^i$ and $\tilde{\pi}_{(-i)}^j$.

We define the gains from trade as $GFT_{ij}^k \equiv \pi_{ij}^k - \tilde{\pi}_{(-)}^k$ for $k \in \{i, j\}$, which measure the value each firm derives from the match. For the exporter, these gains correspond to revenue from j net of the incremental cost of serving that buyer. For the importer, losing variety i raises the input price index p_j^f , increasing marginal costs, reducing optimal output, and lowering downstream revenues. In our setting, these gains are strictly positive.¹¹

The negotiated price solves a Nash bargaining problem that splits these gains, with $\phi \in (0, 1)$ denoting the importer's exogenous bargaining weight:

$$\max_{p_{ij}, q_{ij}} (GFT_{ij}^i)^{1-\phi} (GFT_{ij}^j)^\phi \quad \text{s.t.} \quad q_{ij} = q_j^f \varsigma_{ij}^\rho \left(\frac{p_{ij}}{p_j^f} \right)^{-\rho}. \quad (2.7)$$

We adopt the Nash-in-Nash solution concept, taking prices and quantities on all other links

¹⁰This protocol is analogous to the employer “right-to-manage” model in labor economics (e.g., [Azkarate-Askasua and Zerecero, 2025](#)), where wages are bargained while firms retain the choice of employment.

¹¹Appendix B.3 discusses an extension with more general nonparametric outside options.

as given (Horn and Wolinsky, 1988; Collard-Wexler et al., 2019).

2.3 Price Characterization

To characterize equilibrium prices, it is useful to summarize each bilateral relationship with two market shares that capture each party's relative importance to the other. For a given link (i, j) , the exporter's *supplier share* s_{ij} is its share of the importer's foreign input expenditure; the importer's *buyer share* x_{ij} measures its share of the exporter's total output. Formally,

$$s_{ij} \equiv \frac{p_{ij}q_{ij}}{\sum_{k \in \mathcal{Z}_j} p_{kj}q_{kj}}, \quad \text{and} \quad x_{ij} \equiv \frac{q_{ij}}{\sum_{k \in \mathcal{Z}_i} q_{ik}}. \quad (2.8)$$

We also define $\xi(GFT_{ij}^j, p_{ij})$ as the elasticity of the importer's gains from trade with respect to the bilateral price:

$$\xi(GFT_{ij}^j, p_{ij}) \equiv -\frac{d \ln GFT_{ij}^j}{d \ln p_{ij}} = \frac{(\eta - 1)s_{ij}}{1 - (1 - s_{ij})^{\frac{\eta-1}{\rho-1}}}. \quad (2.9)$$

With these objects in hand, we can now present our first key theoretical result.

Proposition 1. *The solution to (2.7) yields the price-cost ratio $\mu_{ij} \equiv p_{ij}/MC_i$, which satisfies*

$$\mu_{ij} = (1 - \omega_{ij}) \mu^{\text{oligopoly}}(s_{ij}) + \omega_{ij} \mu^{\text{oligopsony}}(x_{ij}), \quad (2.10)$$

with

$$\omega_{ij} = \frac{\phi \xi(GFT_{ij}^j, p_{ij})}{\phi \xi(GFT_{ij}^j, p_{ij}) + (1 - \phi)(\varepsilon_{ij} - 1)} \in (0, 1). \quad (2.11)$$

Proof. See Appendix A.1.

The first component in equation (2.10), $\mu_{ij}^{\text{oligopoly}}$, is the standard *oligopoly markup* arising under Bertrand competition with differentiated products:

$$\mu_{ij}^{\text{oligopoly}} = \frac{\varepsilon_{ij}}{\varepsilon_{ij} - 1} \geq 1 \quad \text{where} \quad \varepsilon_{ij} = (1 - s_{ij})\rho + s_{ij}\eta. \quad (2.12)$$

This expression coincides with the equilibrium price-cost ratio when the exporter has full bargaining power ($\phi = 0$), in which case $\omega_{ij} = 0$.

Here, the residual demand elasticity ε_{ij} is a weighted average of the within-input elasticity

ρ and the elasticity

$$\eta \equiv -\frac{d \ln q_j^f}{d \ln p_j^f} = \frac{(\varrho - \gamma) + \nu(1 - (\varrho - \gamma))}{\varrho + \nu(1 - \varrho)}, \quad (2.13)$$

which depends on technology (γ, ϱ) and demand (ν) .¹² When $\rho > \eta$, the markup increases in the supplier share s_{ij} , reflecting greater oligopoly power.

The second component in equation (2.10), $\mu_{ij}^{\text{oligopsony}}$, is given by

$$\mu_{ij}^{\text{oligopsony}} = \theta \cdot \frac{1 - (1 - x_{ij})^{1/\theta}}{x_{ij}} \leq 1. \quad (2.14)$$

This expression corresponds to the equilibrium price–cost ratio when the importer has full bargaining power ($\phi = 1$), in which case $\omega_{ij} = 1$. Since $\mu_{ij}^{\text{oligopsony}} \leq 1$, we refer to it as an *oligopsony markdown*.¹³

The oligopsony markdown has a natural quasi-rents interpretation. Under full importer bargaining power, the importer extracts all exporters’ surplus and bargain a price equals the exporter’s incremental cost of serving them. Expressed relative to current marginal cost, the corresponding price–cost ratio is

$$\mu_{ij}^{\text{oligopsony}} = \frac{TC_i(q_i) - TC_i(q_i - q_{ij})}{MC_i(q_i) q_{ij}}.$$

The numerator is the exporter’s actual cost saving if the relationship with buyer j is severed. The denominator is what that cost would be if all units sold to j were valued at current marginal cost. Under decreasing returns, the incremental cost of serving j lies below current marginal cost, allowing the negotiated price to fall below marginal cost. This gap, which owes to quasi rents, captures oligopsony power in the model.


This markdown decreases with the buyer share x_{ij} . Larger buyers account for a greater share of the exporter’s output, so walking away from the relationship generates larger cost savings for the exporter and strengthens the importer’s bargaining position. As $x_{ij} \rightarrow 0$, the buyer is atomistic, the incremental cost of serving the match converges to marginal cost, and $\mu_{ij}^{\text{oligopsony}} \rightarrow 1$, so that $p_{ij} = MC_i$. As $x_{ij} \rightarrow 1$, the importer is a monopsonist, the negotiated price converges to average cost, and $\mu_{ij}^{\text{oligopsony}} \rightarrow AC_i/MC_i = \theta$.¹⁴

¹²This is different than in standard nested-CES models with Bertrand competition, where η is typically a primitive preference parameter (e.g., [Atkeson and Burstein, 2008](#)).

¹³Although we refer to $\mu_{ij}^{\text{oligopsony}}$ as a markdown, it is formally a price–cost ratio. This differs from the conventional definition of a markdown as the wedge between the input price and the input’s marginal revenue product.

¹⁴Our oligopsony mechanism is observationally equivalent to quantity discounts across a supplier’s buyers:

FIGURE 1: Equilibrium Allocations With Different Bargaining Power



Notes: Each panel illustrates the negotiated price and quantity under a different value of ϕ . The exporter's residual demand function $p(q_{ij})$ is shown in black; the importer's residual supply function $MC_i(q_{ij})$ is shown in blue and the average cost function $AC_i(q_{ij})$ is in red. All panels assume $\theta < 1$.

This mechanism operates only under decreasing returns to scale. When $\theta = 1$, marginal and average costs coincide, leaving no surplus to extract. As a result, the negotiated price cannot fall below marginal cost, and the markdown in (2.14) equals one for all x_{ij} .

Bargaining Weights. Equation (2.10) expresses the bilateral markup as a weighted average of the oligopoly markup (2.12) and the oligopsony markdown (2.14). The weight ω_{ij} governs the relative importance of these two forces. It is increasing in the exogenous bargaining parameter ϕ , but it is also endogenous because it depends on both the importer's gains-from-trade elasticity $\xi(GFT_{ij}^j, p_{ij})$ and the residual demand elasticity ε_{ij} . Thus, even for a given value of ϕ , effective bargaining power varies systematically across matches.

In particular, for $\phi \in (0, 1)$, the weight ω_{ij} is hump-shaped in the supplier share s_{ij} . When s_{ij} is small, the bilateral price has little effect on the importer's overall costs, so the importer has little to gain from bargaining aggressively. When s_{ij} is very large, the importer is highly exposed to the relationship but also depends heavily on it, which weakens its outside option. Effective bargaining power is therefore strongest at intermediate supplier shares, where cost exposure is significant, but substitution remains feasible.

Equilibrium Allocation. Figure 1 illustrates the equilibrium allocation for a single i - j match under alternative bargaining weights ϕ . The exporter faces a downward-sloping residual demand curve and its associated marginal revenue curve, shown in black and grey, re-

conditional on a supplier and product, marginal cost is common across buyers, yet buyers that account for a larger share of the supplier's output (x_{ij}) obtain higher markdowns and pay lower negotiated prices.

spectively. The importer faces an upward-sloping residual marginal and average cost curves, shown in blue and red.

Panel (A) shows the oligopoly benchmark ($\phi = 0$). Quantity is pinned down by the intersection of marginal revenue and marginal cost, and price is then determined from residual demand. Exporter rents can be decomposed into two components: oligopoly rents, shown in dark blue, equal to $(p - MC)q$, and quasi-rents, shown in light blue, equal to the gap between marginal and average cost under decreasing returns, $(MC - AC)q$.

Panel (B) shows full importer bargaining power ($\phi = 1$). The negotiated allocation lies at the intersection of residual demand and average cost curves, leaving the exporter with zero net gains from trade. In this case, the importer captures the full quasi-rents generated by decreasing returns. Relative to the oligopoly benchmark, the negotiated price is lower and the traded quantity is higher.

Panel (C) shows the intermediate case ($\phi \in (0, 1)$). The negotiated outcome lies between the two polar regimes, consistent with Proposition 1. The striped area highlights the corresponding range of price–marginal cost ratios, which may lie either above or below one depending on bargaining power and the magnitude of quasi-rents. The grey shaded region represents total exporter surplus, equal to $(p - AC)q$, as bargaining power varies between the two extremes. As ϕ rises, the importer captures a larger share of the surplus, so exporter rents decline and converge to zero at $\phi = 1$.

Markups and Market Shares. Both components of the bilateral markup vary systematically with bilateral market shares. A higher supplier share s_{ij} raises the oligopoly markup, while a higher buyer share x_{ij} lowers the oligopsony markdown. The co-movements between prices and bilateral shares are thus informative about the presence of the two forces, as summarized in the following proposition.

Proposition 2. The bilateral markup μ_{ij} satisfies:

1. **Supplier Share.** If μ_{ij} increases in s_{ij} , then $\phi < 1$.
2. **Buyer Share.** If μ_{ij} decreases in x_{ij} , then $\phi > 0$ and $\theta < 1$.

Proof: See Appendix A.2.

These cross-sectional predictions guide the empirical analysis in Section 3.4.

Clarification of Terminology We use *buyer power* or *importer market power* (equivalently, *countervailing power*) to refer broadly to the importer’s ability to influence prices

through bargaining (i.e., any $\phi > 0$). In contrast, we reserve *oligopsony power* for the case in which a higher buyer share x_{ij} allows the importer to negotiate lower prices, other things equal. Oligopsony power therefore requires both $\phi > 0$ and $\theta < 1$, as it relies on the presence of quasi-rents in the exporter's technology.

2.4 Pass-Through Characterization

Our pricing framework allows for an analytical characterization of the short-run effects of unanticipated tariffs. Let T_c denote the gross tariff rate. We assume that the tariff enters multiplicatively into the price charged by exporter i from country c to importer j :

$$\ln p_{ij} = \ln \mu_{ij} + \ln MC_i + \ln T_{c(i)}.$$

While a tariff on country c may affect all relationships involving its exporters, we treat it as a pair-specific shock by holding all other match prices fixed: $dp_{rj} = 0$ for all $r \in \mathcal{Z}_j \setminus i$ and $dp_{i\ell} = 0$ for all $\ell \in \mathcal{Z}_i \setminus j$.

The following proposition characterizes this direct component of tariff pass-through.

Proposition 3. *The tariff pass-through elasticity into the bilateral import price p_{ij} , holding fixed all other prices in the network and general equilibrium variables, is:*

$$\Phi_{ij} \equiv \frac{d \ln p_{ij}}{d \ln T_{c(i)}} = \frac{1}{1 + \Gamma_{ij} + \Lambda_{ij}}, \quad (2.15)$$

where:

$$\Gamma_{ij} \equiv -\frac{d \ln \mu_{ij}}{d \ln p_{ij}} \quad \text{and} \quad \Lambda_{ij} \equiv -\frac{d \ln MC_i}{d \ln p_{ij}}$$

denote the partial elasticities of the equilibrium markup μ_{ij} and the exporter's marginal cost MC_i to changes in the bilateral price p_{ij} , respectively.

Proof: See Appendix A.3.

Proposition 3 highlights two distinct mechanisms shaping tariff pass-through in firm-to-firm trade: a markup channel, capturing strategic markup responses, and a cost channel, capturing how the exporter's marginal cost adjusts in response to price changes.¹⁵ We characterize each channel in turn.

¹⁵Appendix B.4 extends equation (2.15) to incorporate indirect effects, for example shocks to exporter i that affect other prices and quantities and in turn p_{ij} . While we do not model these dynamics explicitly, the extension illustrates how spillovers across relationships can drive a wedge between reduced form and structural estimates.

The Markup Channel (Γ_{ij}). The elasticity Γ_{ij} captures strategic markup responses to price changes. It can be written as

$$\Gamma_{ij} = \left[(1 - \omega_{ij}^\Gamma) \Gamma_{ij}^{\text{oligopoly}} + \omega_{ij}^\Gamma \Gamma_{ij}^{\text{oligopsony}} \right] + \left(1 - \frac{\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}} \right) \Gamma_{ij}^\omega,$$

where $\omega_{ij}^\Gamma \equiv \omega_{ij} \frac{\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} \in [0, 1]$, $\Gamma_{ij}^{\text{oligopoly}}$ and $\Gamma_{ij}^{\text{oligopsony}}$ denote the oligopoly and oligopsony markup elasticities, and Γ_{ij}^ω is the elasticity of the bargaining weight with respect to the bilateral price.

The term in square brackets captures the direct markup adjustment, holding bargaining weights (ω_{ij}) fixed. The second term reflects the endogenous response of these weights.

When exporter bargaining power dominates ($\phi \rightarrow 0$), Γ_{ij} converges to $\Gamma_{ij}^{\text{oligopoly}}$, given by

$$\Gamma_{ij}^{\text{oligopoly}} = - \frac{d \ln \mu_{ij}^{\text{oligopoly}}}{d \ln p_{ij}} = \frac{1}{\varepsilon_{ij} - 1} \cdot \frac{\rho - \varepsilon_{ij}}{\varepsilon_{ij}} \cdot (\rho - 1)(1 - s_{ij}) \geq 0.$$

In this regime, the markup channel operates through *strategic complementarities*, as is standard in oligopoly models.¹⁶ An increase in the bilateral price induces the exporter to compress its markup in order to limit trade diversion, generating incomplete pass-through. The response is strongest at intermediate supplier shares, implying a U-shaped relationship between $\Gamma_{ij}^{\text{oligopoly}}$ and s_{ij} .¹⁷

When importer bargaining power dominates ($\phi \rightarrow 1$), Γ_{ij} converges to $\Gamma_{ij}^{\text{oligopsony}}$, given by

$$\Gamma_{ij}^{\text{oligopsony}} = - \frac{d \ln \mu_{ij}^{\text{oligopsony}}}{d \ln p_{ij}} = \left(\frac{x_{ij}(1 - x_{ij})^{\frac{1}{\theta} - 1}}{\theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]} - 1 \right) (1 - x_{ij}) \varepsilon_{ij} \leq 0.$$

In this case, the elasticity Γ_{ij} is governed by markdown adjustments and reflect *strategic substitutabilities* among importers. As p_{ij} rises, importer j contracts demand, reducing quasi-rents and weakening its bargaining position. The resulting markdown compression amplifies the price response and can generate more-than-complete pass-through.

The elasticity $\Gamma_{ij}^{\text{oligopsony}}$ depends on both the buyer share x_{ij} and the supplier share s_{ij} . It is U-shaped in x_{ij} , vanishing as $x_{ij} \rightarrow 0$ (atomistic buyer) or $x_{ij} \rightarrow 1$ (monopsonist), and

¹⁶See, e.g., [Amiti et al., 2014](#); [Auer and Schoenle, 2016](#); [Garetto, 2016](#); [Amiti et al., 2019a](#).

¹⁷[Amiti et al. \(2014\)](#) show that, in their model, a first-order approximation around small supplier shares implies that pass-through declines with exporter share. We cannot rely on this approximation in our bilateral pricing framework, where both low and high shares are observed.

peaking at intermediate values. It declines in s_{ij} because a higher supplier share reduces the demand elasticity ε_{ij} , limiting variation in x_{ij} and weakening the markdown response.

With two-sided market power ($\phi \in (0, 1)$), the elasticity Γ_{ij} reflects a convex combination of oligopoly and oligopsony markup elasticities, with weights ω_{ij}^Γ proportional to importer bargaining power. It also reflects the endogenous response of bargaining weights through $\Gamma_{ij}^\omega \equiv \frac{d \ln \omega_{ij}}{d \ln p_{ij}}$, whose sign and magnitude depend on s_{ij} .¹⁸ As $s_{ij} \rightarrow 0$ or $s_{ij} \rightarrow 1$, ω_{ij} converges to the exogenous parameter ϕ , implying $\Gamma_{ij}^\omega \rightarrow 0$.

The Cost Channel (Λ_{ij}). The cost elasticity Λ_{ij} captures how marginal costs respond to price changes. An increase in prices reduces demand and, under decreasing returns, lowers marginal costs. Formally,

$$\Lambda_{ij} = \frac{1 - \theta}{\theta} x_{ij} \varepsilon_{ij} \geq 0. \quad (2.16)$$

The first term, $(1 - \theta)/\theta$, captures the slope of marginal cost or, equivalently, the *inverse export supply elasticity* at the exporter level. The buyer share x_{ij} captures the exposure of the exporter's total output to the relationship: importers that account for a greater fraction of the exporter's output exert stronger scale effects on marginal cost. The residual demand elasticity ε_{ij} captures how strongly bilateral quantity responds to a change in price. The product $(1 - \theta)/\theta \cdot x_{ij}$ is the *residual inverse export supply elasticity* faced by buyer j : it scales with the buyer share, reflecting each buyer's exposure to the exporter's overall supply curvature. Together, x_{ij} and ε_{ij} determine the strength of the match-specific scale effect.


This mechanism may appear similar to a standard cost channel driven by firm-level decreasing returns. While upward-sloping export supply can generate incomplete pass-through, a model with atomistic buyers attributes this entirely to supply curvature. As shown in Appendix B.5, this implies uniform pricing and identical pass-through across buyers sourcing from the same exporter. In contrast, our framework features non-atomistic buyers and produces relationship-specific residual cost elasticities that scale with buyer shares, generating cross-sectional heterogeneity in pass-through.

Pass-Through Elasticity and Bilateral Market Shares Figure 2 plots the pass-through elasticity Φ_{ij} as a function of s_{ij} and x_{ij} for three values of θ (rows) and ϕ (columns).¹⁹ The figure shows that the same primitives governing oligopsony power also shape pass-through. Variation in θ (vertical) captures how exporters' returns to scale, via the inverse export supply elasticity, determine the balance between cost and markup channels, while variation

¹⁸The exact expression for this term is provided in Appendix A.3.

¹⁹We set $\theta = 0.9$ as an intermediate value of the returns-to-scale parameter, consistent with estimates in the production function literature, typically between 0.9 and 1; see, e.g., Basu and Fernald (1997).

FIGURE 2: Pass-Through Elasticity and Bilateral Market Shares



Notes: The figure presents heatmaps of the pass-through elasticity Φ_{ij} as a function of market shares s_{ij} (x-axis) and x_{ij} (y-axis), for $\phi \in \{0, 0.5, 1\}$ (columns) and $\theta \in \{0.5, 0.9, 1\}$ (rows). Other parameters are fixed at $\gamma = 0.5$, $\varrho = 1$, $\nu = 4$, and $\rho = 10$.

in ϕ (horizontal) captures how bargaining power governs strategic markup adjustment.

Moving from top to bottom, θ rises from 0.5 to 1. At $\theta = 0.5$, export supply is steep, so marginal cost is highly sensitive to output and the cost channel dominates across the (s_{ij}, x_{ij}) space, rendering pass-through largely invariant to ϕ and the markup channel. As θ increases, the cost channel weakens and markup adjustment becomes more important. Markdown responses also attenuate, but the cost channel continues to dominate in x_{ij} (Appendix A.4), so pass-through becomes less incomplete. At $\theta = 1$, supply is flat, the cost channel vanishes, and pass-through is governed entirely by markup adjustment.

Moving left to right, higher ϕ traces the role of bargaining power in shaping strategic markup adjustment. At $\phi = 0$, the markup channel reflects strategic complementarities that dampen

pass-through; at $\phi = 1$, it reflects strategic substitutabilities associated with oligopsony power that amplify pass-through. This gradient is most pronounced in the bottom row ($\theta = 1$), where only the markup channel operates, and attenuates as θ falls. At $\theta = 0.5$, markup adjustment is second order.

A key implication of Figure 2 is that pass-through is incomplete over most of the parameter space, with the sole exception of the constant-returns, full-importer-power case ($\theta = 1$, $\phi = 1$). Different (θ, ϕ) pairs generate incomplete pass-through through distinct mechanisms and to varying degrees. In the lower-left region of the heatmaps—corresponding to the low-share range where most bilateral matches lie in the data (Table 1)—average pass-through is similar across specifications, making models difficult to distinguish on that margin alone. By contrast, the sign and magnitude of the pass-through gradient in x_{ij} differ sharply across variants. The following proposition formalizes this contrast, which we exploit in Section 4.3 to discriminate among models using cross-sectional pass-through moments.

Proposition 4. *When $\theta < 1$, the pass-through elasticity Φ_{ij} decreases with the importer’s buyer share x_{ij} .*

Proof: See Appendix A.4. Intuitively, when $\theta < 1$, the cost channel dominates the mark-down channel. Equation (2.16) implies that pass-through falls with x_{ij} .

The relationship between Φ_{ij} and s_{ij} is, by contrast, ambiguous. Under strategic complementarities, pass-through is U-shaped in s_{ij} . Under strategic substitutabilities and through the cost channel, it is increasing in s_{ij} . The former pattern arises when bargaining power is low and returns to scale are close to constant, while the latter emerges as either force strengthens. Which effect dominates is ultimately an empirical matter.

2.5 Discussion

We briefly review our key modeling choices before turning to the empirical analysis.

Bargaining Protocol and Quantities Our baseline model assumes *demand-driven quantities*, whereby the importer chooses quantities given negotiated prices. This ensures allocative pricing and nests the standard monopolistic-competition benchmark as a limiting case, consistent with firm-level evidence in international trade (e.g., [Gopinath and Itskhoki, 2011](#)).

We derive two alternative protocols in the Appendix. The first is *efficient bargaining* (Appendix B.1), which maximizes joint surplus but implies non-allocative prices, a prediction at odds with the data. The second is *supply-driven bargaining* (Appendix B.2), a monopsony-style benchmark, which implies that prices and quantities move together along the exporter’s

supply curve. This implication is inconsistent with our empirical evidence that tariff-induced price increases are accompanied by quantity declines.

One caveat of our framework is that when ϕ is sufficiently high, equilibrium prices can fall below the efficient level ($p < p^{\text{eff}}$), implying quantities above the competitive benchmark ($q > q^{\text{eff}}$). Recent theories of bilateral oligopoly rule out this case by imposing a supplier participation constraint that prevents quantities from exceeding the exporter’s supply (marginal cost) curve when $p < p^{\text{eff}}$ (e.g., [Demirer and Rubens, 2025](#); [Avignon et al., 2024](#)). We abstract from this extension because, at high ϕ , it places equilibrium on the upward-sloping supply curve, implying positive price-quantity comovement, contrary to our evidence.

General Equilibrium and Network Spillovers Our theoretical results are derived in partial equilibrium, holding wages, aggregate demand shifters, and competitors’ prices and quantities fixed. The pass-through elasticity derived in [Section 2.4](#) therefore isolates the direct, intensive-margin response of markups and marginal costs to a tariff shock.

Embedding this framework in a general equilibrium environment introduces two main complexities. First, under decreasing returns, quantity adjustments along one link affect the exporter’s marginal cost across all links. Second, aggregate shocks or network spillovers may differentially affect outside options across matches of the same exporter. [Appendix B.3](#) discusses how generalizing outside options modifies the bargaining problem. Fully accommodating these channels, however, requires committing to a specific macroeconomic closure for wages, final demand, and network formation, which is beyond the scope of this paper. As we show below, the partial equilibrium approach is supported by the close fit between model predictions and untargeted pass-through moments.

3 Data and Stylized Facts

This section describes the data and preliminary empirical analysis. [Section 3.1](#) outlines the main data sources. [Section 3.2](#) discusses how we adapt the baseline model to the data to construct key variables. [Section 3.3](#) details the sample selection and provides summary statistics. Finally, [Section 3.4](#) presents reduced-form evidence on prices and pass-through through the lens of the model.

3.1 Data Sources

Our main dataset is the U.S. Census Bureau’s Linked/Longitudinal Firm Trade Transaction Database (LFTTD), which covers the universe of U.S. import transactions from 2001 to

2018. Each observation corresponds to a shipment from a foreign exporter to a U.S. importer and includes the transaction date, product classification at the 10-digit Harmonized System (HS10) level, Free on Board (FOB) import value in U.S. dollars, physical quantity, transportation mode, and country of origin. Exporters are identified using a manufacturer ID (MID) constructed by the Census Bureau from the exporter’s name, street address, city, and country.²⁰

To focus on arm’s-length trade, we exclude related-party transactions from the baseline sample. The LFTTD includes a related-party indicator based on a mandatory field in U.S. Customs forms, flagging relationships with ownership stakes of at least five percent. While widely used, this measure may misclassify firms due to its reliance on self-reporting and a low reporting threshold (Ruhl, 2015). To improve accuracy, we construct an alternative indicator using ORBIS, which provides firm-level cross-border ownership links. We merge ORBIS to the LFTTD as described in Appendix C.1.

We supplement the transaction-level data with information on statutory U.S. import tariffs introduced during the 2018 trade war. We use the dataset from Fajgelbaum et al. (2020), which records the timing, product coverage, and country-specific scope of these measures at the HS8-month level. The tariffs averaged 25 percentage points and were imposed on top of existing rates, targeting selected goods. They were implemented in phases over the course of the year, beginning with imports from China and later expanding to goods from other trade partners, including Canada, Mexico, and the European Union. Tariff changes are annualized based on the number of months each measure was in effect.

3.2 Measuring Key Variables of the Model

To construct the key variables of interest, we extend the model to include multiple foreign inputs, indexed by h . Each input corresponds to an HS10 product category. We model the foreign input bundle as a Cobb-Douglas composite of individual product quantities:

$$q_j^f = \prod_{h \in \mathcal{H}_j} \left(q_{jh}^f \right)^{\alpha_{jh}}, \quad \text{where} \quad q_{jh}^f = \left(\sum_{i \in \mathcal{Z}_j^h} s_{ijh} \cdot (q_{ijh})^{\frac{\rho-1}{\rho}} \right)^{\frac{\rho}{\rho-1}},$$

and $\alpha_{jh} \in (0, 1)$ denotes the (observed) Cobb-Douglas share of input h in firm j ’s total imports of foreign intermediates. This formulation implies that the elasticity of the importer’s

²⁰The MID combines country code, elements of the firm name, city, and address (Kamal and Monarch, 2018). Because the algorithm is not standardized, identifiers may be inconsistent: one firm may have multiple MIDs or multiple firms may share one. Following Kamal and Monarch (2018), we construct a robustness measure that truncates location fields to improve consistency. Results are robust.

marginal cost with respect to the price of foreign input h is $\frac{d \ln c_j}{d \ln p_{jh}^f} = \alpha_{jh} \gamma \in (0, 1]$. We also note that allowing this elasticity to vary across importer-product pairs in turn makes η in equation (2.13) to vary across the same dimension, η_{jh} .

We construct the exporter’s supplier share as $s_{ijh} \equiv \frac{p_{ijh} q_{ijh}}{\sum_{k \in \mathcal{Z}_j^h} p_{kjh} q_{kjh}}$, where $\mathcal{Z}_j^h$ denotes the set of firm j ’s foreign suppliers of input h . The numerator captures the total value of imports of product h from exporter i (a MID in our data) to firm j in a given year. The denominator aggregates imports of product h from all foreign suppliers to j .

In contrast, the importer’s buyer share is constructed as $x_{ijh} \equiv \frac{q_{ijh}}{\sum_{k \in \mathcal{Z}_i^h} q_{ikh}}$, where $\mathcal{Z}_i^h$ is the set of all U.S. importers buying product h from exporter i . Since our dataset only includes U.S. importers, we assume that exporter i operates product- and destination-specific production lines. Under this assumption, the denominator of x_{ijh} , which captures the total quantity of product h from exporter i , includes only those sold to U.S. buyers. This restriction reflects a data limitation, as we do not observe importer destinations beyond the U.S. and thus cannot account for the full set of an exporter’s buyers.²¹

3.3 Sample Construction and Summary Statistics

We apply a sequence of restrictions to the LFTTD to construct the estimation sample. Some restrictions are required for identification and counterfactuals; others align the baseline sample with the model’s environment. Appendix C.2 provides details, and Appendix Table C.1 traces the sample under each restriction.

Essential restrictions. Three restrictions are essential. First, we require importer-exporter-product triplets to be observed in two consecutive years, a necessary restriction for bilateral price to be computed. Second, because identification of the bargaining and returns-to-scale parameters relies on cross-sectional price dispersion across importers sourcing the same exporter-product-year, we restrict the sample to exporter-product pairs that transact with at least two U.S. buyers in consecutive years. Third, to limit measurement error and ensure that unit values more closely reflect prices rather than compositional changes within HS10 products, we follow standard practice and (i) exclude energy goods, (ii) drop unit values outside the 1st-99th percentile within product, and (iii) remove observations with absolute log price changes exceeding four.

²¹To address potential overstatement of buyer share x_{ij} due to unobserved exports to other destinations, we replicate the analysis for Canada and Mexico, which send most exports to the U.S. (71% and 73% in 2019). Estimates are similar in this subsample. Full results available upon request.

Other restrictions. We impose two additional restrictions to align the data with the model. First, we exclude related-party transactions, as prices in such relationships may reflect forces, such as transfer pricing, outside our framework (Bernard et al., 2006). In the baseline, pairs are classified as related if ORBIS identifies a common corporate parent; alternative definitions based on the LFTTD related-party flag or multinational status yield similar results.

Second, we focus on capital and intermediate inputs by excluding consumption goods under the Broad Economic Categories (BEC) classification. This reflects the model’s focus on supply chains, where lock-in from relationship-specific investments limits short-run scalability and is more pronounced upstream (e.g., Nunn, 2007). In robustness exercises, however, we show that the main results extend to the full set of imported goods.

After all restrictions, the baseline sample covers approximately \$160 billion in import value and 250,000 buyer-supplier-product-year observations in 2017-2018. This corresponds to about 20% of total U.S. imports by value and roughly 15% of buyer-supplier-product triplets. Although not representative of aggregate U.S. trade, the sample isolates repeated, arm’s-length production relationships where relationship-level price changes can be credibly measured and the model’s mechanisms credibly identified.

Summary Statistics Table 1 reports summary statistics for our final sample.²² Panel A documents substantial concentration on both sides of the market. On average, an exporter supplies 32% of an importer’s total imports of a given HS10 product, with a median share of 15%. The average buyer share is lower, at 25%, with a median of 10%, likely reflecting the exclusion of cases with buyer share equal to one. The two shares are highly dispersed and essentially uncorrelated (correlation 0.04).

Trading relationships are also persistent: the typical importer-exporter pair has been transacting in a given HS10 product for about four years, and in any product for nearly five. Activity within the relationship is intensive: at the median, a pair has accumulated 16 transactions in a product to date and trades 2.5 distinct HS10 products in a given year. Buyer and supplier tenure across all partners and products average about seven and six years, respectively.

Panel B of Table 1 reports descriptive statistics for three bilateral (log) price measures, each defined as transaction value divided by quantity: an FOB (Free on Board) pre-duty price that excludes shipping and insurance charges; a CIF (Cost, Insurance, and Freight) pre-duty price that adds those charges; and a post-duty price that further incorporates import

²²Appendix Table C.2 provides additional statistics. Appendix Table C.3 and Table C.4 report sample composition and summary statistics for the pooled sample including all BEC categories.

TABLE 1: Summary Statistics for Main Estimation Sample (2001–2018)

Variable	Mean	Median	P25–P75
<i>Panel A: Trade Relationships</i>			
s_{ijh} : Supplier share	0.32	0.15	0.03–0.57
x_{ijh} : Buyer share	0.25	0.10	0.02–0.40
Relationship length	4.0	3.5	2.5–5.5
Relationship length (<i>all products</i>)	4.8	4.5	2.5–6.5
# Transactions	120	16	6.5–50
# Products per pair	3.8	2.5	1.5–4.5
Buyer tenure	6.9	6.5	3.5–10.0
Supplier tenure	6.4	6.5	3.5–9.5
<i>Panel B: Prices (log unit values)</i>			
$\log p$ (pre-duty, excl. charges)	3.40	3.00	1.30–5.40
$\log p$ (pre-duty)	3.50	3.10	1.40–5.40
$\log p^{\text{duty}}$ (post-duty)	3.50	3.10	1.40–5.40

Notes: This table reports summary statistics for the main estimation sample at the importer-supplier-product level. Full details on sample construction are provided in Appendix C.2. Variables are measured at the HS10 product level unless noted otherwise. Panel A reports trade relationship characteristics. s_{ijh} is exporter i 's share in buyer j 's imports of product h ; x_{ijh} is buyer j 's share in exporter i 's U.S. exports of product h . Relationship length is the number of years between the supplier-buyer pair's first transaction in product h and year t ; Relationship length (*all products*) is defined analogously, using the pair's first transaction in any product. # Transactions is the cumulative number of transactions between the pair in product h , from their first transaction through year t . # Products per pair is the number of distinct HS10 products traded by the pair in year t . Buyer tenure (Supplier tenure) is the number of years between the buyer's (supplier's) first recorded transaction in the LFTTD and year t , pooling across partners and products. Panel B reports log unit values (value divided by quantity). $\log p$ (pre-duty, excl. charges) corresponds to the FOB (Free on Board) unit value, excluding shipping and insurance charges; $\log p$ (pre-duty) includes these charges to obtain the CIF (Cost, Insurance, and Freight) unit value; and $\log p^{\text{duty}}$ (post-duty) further incorporates import duties. Statistics are based on confidential LFTTD data and rounded according to U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

duties. All three are highly dispersed across importer-exporter-product-year observations, with interquartile ranges of about four log points.

To assess the sources of this variation, we conduct a variance decomposition following [Fontaine et al. \(2020\)](#) (Appendix D.1). Product-year fixed effects account for about 50% of the total variance, while match-specific residuals account for 4%. Focusing on within supplier-product-year variation, 77% of the remaining dispersion is explained by match-specific factors. This pattern holds across price definitions and underscores the importance

of relationship-specific forces in pricing.

3.4 Reduced-Form Evidence on Prices and Pass-Through

As a preliminary exploration of mechanisms, we examine the empirical relationships between markups, pass-through, and bilateral market shares. We interpret these patterns through the lens of Propositions 2 and 4 as descriptive evidence consistent with two-sided market power and decreasing returns to scale, before turning to structural estimation in Section 4.

Markups and Bilateral Market Shares Proposition 2 implies that under two-sided market power ($\phi \in (0, 1)$) and decreasing returns to scale ($\theta < 1$), bilateral markups rise in the exporter’s supplier share (s_{ijht}) and fall in the importer’s buyer share (x_{ijht}). To characterize these empirical patterns, we estimate

$$\ln p_{ijht} = \alpha_s s_{ijht} + \alpha_x x_{ijht} + \mathbf{X}_{ijht}\boldsymbol{\gamma} + \mathbf{FE} + v_{ijht}, \quad (3.1)$$

where α_s and α_x are expected to be positive and negative, respectively. Because markups are not observed, we use log prices as the dependent variable, and include supplier–product–time fixed effects to absorb cost variation and isolate the markup component.

A concern is that market shares mechanically depend on prices, potentially inducing spurious correlation. To address this, we also report IV estimates using leave-one-out measures of market structure. We instrument s_{ijht} with the average supplier share (excluding i) across other buyers of exporter i (excluding j), and instrument x_{ijht} with the average buyer share (excluding j) across other suppliers to importer j (excluding i).

Table 2 reports the results. Columns (1)-(2) include exporter, importer, and product-year fixed effects. Columns (3)-(4) replace exporter fixed effects with exporter-product-year fixed effects to absorb marginal costs. Columns (5)-(6) further include importer-product-year fixed effects to capture buyer-specific demand shocks. All specifications control for relationship duration (years since the first shipment of product h between i and j). Odd columns report OLS; even columns report IV estimates.

Across specifications, we find $\alpha_s > 0$ and $\alpha_x < 0$, both statistically and economically significant. While the magnitudes are not directly interpretable in isolation, these patterns provide clear support for the model’s core mechanisms. Through the lens of Proposition 2, they are consistent with markup variation reflecting both oligopoly and oligopsony forces.

TABLE 2: Prices and Bilateral Concentration

Dependent Variable:	$\ln p_{ijht}$					
	(1)	(2)	(3)	(4)	(5)	(6)
	OLS	IV	OLS	IV	OLS	IV
s_{ijht}	0.166 (0.003)	0.170 (0.029)	0.237 (0.004)	0.169 (0.029)	0.250 (0.006)	0.642 (0.015)
x_{ijht}	-0.559 (0.003)	-0.093 (0.0174)	-0.617 (0.003)	-0.237 (0.023)	-0.549 (0.006)	-0.783 (0.011)
$FE_i + FE_j + FE_{ht}$	Yes	Yes	No	No	No	No
$FE_{iht} + FE_j$	No	No	Yes	Yes	No	No
$FE_{iht} + FE_{jht}$	No	No	No	No	Yes	Yes
Observations	1,200,000	1,200,000	1,200,000	1,200,000	1,200,000	1,200,000
R-squared	0.958	0.012	0.976	0.041	0.991	0.038
First-stage F stat.	—	5,270	—	3,485	—	19,760
SW F stat (s_{ijht})	—	10,710	—	7,464	—	39,830
SW F stat (x_{ijht})	—	21,120	—	7,197	—	43,550

Notes: This table reports OLS and IV estimates of equation (3.1), where the dependent variable is the log CIF duty-exclusive unit value of product h imported by buyer j from supplier i in year t . Columns alternate between OLS and IV specifications. All regressions control for log relationship length (in years) within HS10 products. Columns (1)–(2) include buyer (FE_j), supplier (FE_i), and product-year (FE_{ht}) fixed effects. Columns (3)–(4) use supplier–product–year (FE_{iht}) and buyer (FE_j) fixed effects. Columns (5)–(6) include fully interacted buyer–product–year and supplier–product–year fixed effects (FE_{jht} , FE_{iht}). IV estimates (even-numbered columns) use leave-one-out instruments: s_{ijht} is instrumented with the average share of other suppliers across buyers of i (excluding j), and x_{ijht} with the average share of other buyers across suppliers to j (excluding i). Because the model includes multiple endogenous regressors, we report both first-stage and conditional F-statistics from Sanderson and Windmeijer (2016). Standard errors are robust. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

Pass-Through and Decreasing Returns to Scale Proposition 4 establishes that, in our framework, a negative comovement between pass-through and buyer share arises if and only if production exhibits decreasing returns to scale ($\theta < 1$).

We test this implication using the reduced-form regression

$$\begin{aligned} \Delta \ln p_{ijht}^* = & \alpha_0 + \alpha_1 \Delta \ln(1 + \tau_{cht}) + \alpha_s \Delta \ln(1 + \tau_{cht}) \cdot s_{ijh,t-1} + \alpha_x \Delta \ln(1 + \tau_{cht}) \cdot x_{ijh,t-1} \\ & + \alpha_2 s_{ijh,t-1} + \alpha_3 x_{ijh,t-1} + \mathbf{X}_{ijht} \boldsymbol{\gamma} + \mathbf{FE} + \epsilon_{ijht}, \end{aligned} \quad (3.2)$$

The dependent variable is the change in the duty-exclusive price, denoted p_{ijht}^* . While the model is expressed in terms of *statutory* tariffs, observed duty-inclusive prices reflect *applied*

duties and may therefore be subject to measurement error. Focusing on duty-exclusive prices isolates price adjustments net of duties and corresponds to the exporter’s received price in the model.²³ We focus on the 2017-2018 period, when U.S. imports were subject to large and unexpected tariff increases.

The interaction terms capture heterogeneity in pass-through with respect to bilateral market shares, $s_{ijh,t-1}$ and $x_{ijh,t-1}$, measured at the beginning of the period. The vector $\mathbf{X}_{ijht}$ includes controls for changes in exporter i ’s sales to other U.S. buyers and the average price change faced by importer j from alternative suppliers, helping to isolate bilateral responses emphasized in the model.

We estimate two specifications. The first includes product-time (FE_{ht}) and exporting country-sector (FE_{cs}) fixed effects, following standard practice in the pass-through literature. This is our baseline. The second is more demanding, combining product-time (FE_{ht}) with importer-time (FE_{jt}) and exporting country-time (FE_{ct}) fixed effects.

Table 3 presents the results using the baseline fixed effects. Column (1) shows that, on average, pass-through into duty-exclusive prices is incomplete: a 10% tariff increase reduces exporter prices by 1.3% , corresponding to 87% pass-through rate. Column (2) shows that this finding is robust to controlling for relationship age. It also reveals that pass-through increases with the length of the relationship, consistent with Heise (2024).

Columns (3), (5), and (6) show that the coefficient on supplier share (α_s) is positive but statistically insignificant, suggesting limited heterogeneity in pass-through on the exporter side. In contrast, the coefficient on buyer share (α_x) (Columns (4)-(6)) is negative and highly statistically significant, indicating that importers with greater buyer share experience lower pass-through. Table D.2 in the Appendix D.2 confirms that these patterns persist under more demanding fixed effects, while Table D.3 shows they are robust to alternative price definitions and the inclusion of general equilibrium controls.

To assess for nonlinearities, we interact tariff changes with quartiles of supplier and buyer shares. Figure D.1 shows no systematic pattern across supplier share quartiles, but a clear, monotonic decline across buyer share quartiles, regardless of the fixed effects.

Taken together, these results indicate that pass-through is largely insensitive to supplier concentration but declines sharply with buyer concentration. By Proposition 4, the negative gradient in x implies $\theta < 1$; these patterns are consistent with the first row of Figure 2.

²³We write the buyer’s effective price as $p_{ijh} = p_{ijh}^*(1 + \tau)$, where p_{ijh}^* denotes the exporter’s received price and τ the statutory tariff rate.

TABLE 3: Pass-Through and Relationship Heterogeneity

Dependent variable:	$\Delta \ln p_{ijht}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1 + \tau_{cht})$	-0.126 (0.091)	-0.168 (0.105)	-0.138 (0.094)	-0.019 (0.097)	-0.032 (0.091)	-0.094 (0.103)
$\Delta \ln(1 + \tau_{cht}) \cdot \ln \text{longevity}_{ijht}$		0.030 (0.022)				0.046 (0.019)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$			0.029 (0.084)		0.033 (0.080)	0.032 (0.079)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$				-0.407 (0.112)	-0.407 (0.112)	-0.416 (0.112)
$\text{FE}_{ht} + \text{FE}_{cs}$	Yes	Yes	Yes	Yes	Yes	Yes
Observations	249,000	249,000	249,000	249,000	249,000	249,000
R-squared	0.04	0.04	0.04	0.04	0.04	0.04

Notes: This table reports estimates of the pass-through of statutory tariffs, $\Delta \ln(1 + \tau_{cht})$, to CIF duty-exclusive prices at the exporter-importer-product-year level, $\Delta \ln p_{ijht}$. Columns (2) and (6) interact tariffs with the log relationship length (in years) within a product h . Columns (3) and (5) interact tariffs with the lagged supplier share, s_{ijht-1} , defined as supplier i 's share in buyer j 's imports of product h . Columns (4) and (5) interact tariffs with the lagged buyer share, x_{ijht-1} , defined as buyer j 's share in supplier i 's exports of product h . All regressions include product-year and exporter country-sector fixed effects ($\text{FE}_{ht} + \text{FE}_{cs}$). Controls include: (i) $\ln \text{longevity}_{ijht}$; (ii) $\Delta \ln q_{i(-j)ht}$, the change in exporter i 's total sales of h to U.S. buyers other than j ; and (iii) $\Delta \ln p_{(-i)jht}$, the weighted average price change charged by other suppliers of h to buyer j , using lagged shares as weights. Standard errors are clustered at the HS8 product and exporter-country level. The sample corresponds to the " + Supplier Multi-Buyer" definition in Table C.1. Observation counts are rounded to four significant digits per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

4 Structural Evidence

The reduced-form patterns in Section 3.4 are consistent with two-sided market power and decreasing returns, but they cannot separately identify the structural parameters or convincingly isolate the model's mechanisms from general-equilibrium effects. We therefore adopt a structural approach. Estimation focuses on ϕ and θ , the two primitives that govern bargaining and decreasing returns; the remaining parameters are calibrated from the literature.

We set the elasticity of substitution across foreign varieties to $\rho = 10$, consistent with Anderson and van Wincoop (2004) and Edmond et al. (2023), who use similar values to match observed U.S. markups. The downstream demand elasticity is set to $\nu = 4$, based on

variety-level import demand estimates in [Broda and Weinstein \(2006\)](#).²⁴

The baseline model features a single foreign input, whereas firms in the data import many (Section 3.2). In the empirical implementation, we therefore replace the common elasticity γ with the term $\alpha_{jh}\gamma$, where α_{jh} is the observed expenditure share of input h in firm j 's imported intermediates. We set $\gamma = 0.5$ equal to the share of imported inputs in total material costs for U.S. manufacturers from [Eldridge and Powers \(2018\)](#).

Finally, the importer's returns-to-scale parameter ϱ has no direct counterpart in the supply-chain production literature. We therefore adopt $\varrho = 1$ as a parsimonious benchmark and assess robustness to alternative values, including $\varrho = 0.5$. Because equilibrium outcomes depend on importer technology (ν, γ, ϱ) only through the composite elasticity η_{jh} , these exercises are equivalent to varying η_{jh} directly.

4.1 Identification and Estimation of the Parameters θ and ϕ

Let Ω_{ijt} denote the information set available to pair (i, j) at the time of negotiation. This includes observed market shares and calibrated parameters. As shown in equation (2.10), conditional on Ω_{ijt} the bilateral markup depends only on the structural parameters (ϕ, θ) , so that $\mu_{ij} = \mu(\phi, \theta \mid \Omega_{ijt})$.

We write the log price of input h sold by exporter i to importer j in year t as

$$\ln p_{ijht} = \ln \mu(\phi, \theta \mid \Omega_{ijt}) + \ln MC_{ijht} \quad \text{with} \quad \ln MC_{ijht} = \frac{1-\theta}{\theta} \ln q_{iht} + \ln k_{ijht},$$

where p_{ijht} is the CIF duty-exclusive unit value and k_{ijht} denotes a match-year-specific cost component.

In the theoretical model, the marginal cost term $(\ln MC_{ijht})$ is common across buyers of a given exporter-product-year. In the empirical implementation, we allow it to vary flexibly at the match-year level to absorb any (i, j, h, t) -specific cost differences in levels. The scale term $\frac{1-\theta}{\theta} \ln q_{iht}$ remains common across buyers within a supplier-product-year and captures the effect of decreasing returns.

Consider exporter i matched with two importers, j and ℓ . Conditional on the joint information set $\Omega_{ij\ell t} \equiv (\Omega_{ijt}, \Omega_{i\ell t})$, we assume mean independence of the unobserved cost component across buyers:

$$\mathbb{E}_k[\ln k_{ijht} - \ln k_{i\ell ht} \mid \Omega_{ij\ell t}] = 0. \quad (4.1)$$

²⁴Appendix E.1 provides additional details.

Differencing log prices across buyers served by the same exporter yields the moment condition

$$\mathbb{E}_k \left[\ln p_{ijht} - \ln p_{ilht} - \left(\ln \mu(\phi, \theta \mid \Omega_{ijt}) - \ln \mu(\phi, \theta \mid \Omega_{ilt}) \right) \mid \Omega_{ij\ell t} \right] = 0. \quad (4.2)$$

After differencing, exporter-level marginal cost components cancel out, so the remaining price differences across buyers reflect differences in bilateral markups. Since the markup function depends only on (ϕ, θ) conditional on Ω_{ijt} , variation in market shares identifies the structural parameters.

Identification requires that equation (4.2) holds uniquely at the true (ϕ, θ) . Because the oligopoly component of the markup does not depend on θ , identification of θ requires $\phi > 0$, so that the oligopsony channel affects pricing. The empirical evidence in Tables 2 and 3 supports the presence of such bilateral effects.

For $\phi \in (0, 1)$, the markup function is strictly monotonic in both parameters over the empirically relevant range. This monotonicity implies local invertibility of the moment condition in (ϕ, θ) . Identification is therefore secured by observing multiple matches for the same exporter, either within a year or over time, under the assumption that bargaining weights are constant across matches.²⁵

Threats to Identification A first concern is endogenous match formation or omitted match characteristics. If such factors affect $\Delta k_{ij\ell ht}$, assumption (4.1) may be violated, potentially biasing estimates of $\hat{\phi}$ and $\hat{\theta}$. Our strategy mitigates this concern by differencing across buyers of the same exporter, which removes shocks common to all matches of that supplier. To further address residual concerns, we implement an instrumental variable strategy; see Appendix E.2 for details.

A second concern is that price dispersion within supplier-product-year cells, even within narrowly defined HS10 categories, may reflect product or service differences rather than bargaining. In the model, such heterogeneity is captured by the idiosyncratic cost component k_{ijht} , which shifts price levels but does not drive identification. Identification instead comes from the systematic relationship between prices and bilateral market shares. To overturn our interpretation, unobserved match characteristics would need to replicate the nonlinear price-share relationship generated by bilateral bargaining under decreasing returns.

Finally, identification does not rely on strong assumptions about the exogeneity of tariffs or other aggregate shocks. Although general equilibrium forces may influence the estimated

²⁵Identification follows from the nonlinearity of equation (2.10) in s_{ijht} and x_{ijht} , which ensures that the derivatives with respect to (ϕ, θ) satisfy a rank condition.

parameters, identification is based solely on cross-sectional variation across buyers within supplier-product-year cells. We therefore require only that such aggregate forces do not differentially affect prices across matches of the same supplier.

Estimation We estimate equation (4.2) via generalized method of moments (GMM),

$$\min_{\{\phi, \theta\}} \mathbf{g}(\phi, \theta)' \mathbf{Z}' \mathbf{W} \mathbf{Z} \mathbf{g}(\phi, \theta), \quad (4.3)$$

where $\mathbf{g}(\phi, \theta)$ stacks all moment conditions in equation (4.2) across all i - j - ℓ pairs and years and $\mathbf{W}$ is the optimal weighting matrix.

To address endogeneity concerns, we first include fixed effects by demeaning $\mathbf{g}(\phi, \theta)$ at the HS10 product, year, and buyer level. This removes average variation across those dimensions, so that only time-varying, pair-specific shocks could bias $\Delta k_{ij\ell ht}$. In addition, we employ instrumental variables ($\mathbf{Z}$) that are plausibly exogenous with respect to the network formation process and other omitted variables.

In particular, the vector $\mathbf{Z}$ includes the total number of importers and exporters in each HS10 product-year, which we interpret as proxies for the pool of potential US buyers and foreign suppliers in a given variety. We also include in $\mathbf{Z}$ the mean and median of the distributions of the two bilateral shares within each year, excluding the focal pairs $i - j$ and $i - \ell$ to preserve over-identification. These instruments vary with the competitive structure within each HS10 product-year and are correlated with the endogenous variables through market structure, but, by construction, are not correlated with the idiosyncratic shocks affecting individual matches.

Extension: pair-specific bargaining weights While our baseline assumes a constant bargaining weight ϕ across all importer-exporter pairs, we also consider an extension to allow ϕ to vary at the pair level.

Given the large number of trade pairs in the data, estimating a separate ϕ_{ij} for each is computationally burdensome. Moreover, our identification strategy does not allow bargaining weights to vary both across pairs and over time. We therefore model bargaining power as a function of observable characteristics:

$$\phi_{ijt} = \frac{\exp(\mathbf{X}_{ijht} \boldsymbol{\kappa})}{1 + \exp(\mathbf{X}_{ijht} \boldsymbol{\kappa})} \in [0, 1], \quad (4.4)$$

where $\boldsymbol{\kappa}$ is a parameter vector to be estimated and $\mathbf{X}_{ijht}$ includes covariates that plausibly influence bargaining outcomes but are not direct determinants of gains from trade in our

model. Specifically, we include: (i) the longevity of the $i - j$ relationship, (ii) the number of transactions between $i - j$ in a year, (iii) the supplier’s outside option relative to the importer’s, measured by the ratio of the quantity of the exporter i ’s sales to buyers other than j in year $t - 1$ over the quantity of the importer j ’ purchases from suppliers other than i in year $t - 1$, and (iv) an indicator variable of whether the buyer and supplier transact multiple HS10 products.

4.2 Estimation Results

We estimate equation (4.3) using data from 2001 to 2016. We exclude 2017 and 2018, which we reserve for out-of-sample validation in Section 4.3 using the tariff shocks during that period. To avoid convergence issues when ϕ is close to one, we estimate the transformed parameter $\bar{\phi} \equiv \ln \frac{\phi}{1-\phi}$, which enters the markup equation linearly.

Table 4 presents the estimation results. Panel B reports the GMM estimates. Columns (1) and (3) assume a constant ϕ , while Columns (2) and (4) allow ϕ_{ij} to vary by trade pair as specified in equation (4.4). The specifications in Columns (1) and (2) are estimated without fixed effects; those in Columns (3) and (4) include year, product, and importer fixed effects. Panel C shows the implied values of ϕ or ϕ_{ijt} .

The parameters are precisely estimated. Across specifications, U.S. importers appear to wield substantial bargaining power, with estimated values of ϕ ranging from 0.7 to 0.9. Our preferred estimate, reported in Column (1), is $\hat{\phi} = 0.83$, implying that U.S. importers have, on average, roughly four times the bargaining power of their foreign suppliers.

The returns to scale parameter $\hat{\theta}$ is consistently estimated below one, ranging from 0.42 to 0.51 across specifications, with a preferred estimate of 0.46.²⁶ This implies an inverse residual export supply elasticity of about 0.24-0.35 for the average importer.²⁷ These values are consistent with evidence from U.S. manufacturing under short-run constraints: [Boehm and Pandalai-Nayar \(2022\)](#) report median inverse elasticities around 0.3 at typical capacity levels, and [Broda et al. \(2008\)](#) document similarly low elasticities across many traded goods.

Moving to the estimates of the vector $\hat{\kappa}$, the results are mixed across the relationship-level controls. Relationship longevity is significant in both specifications but changes sign when fixed effects are included, suggesting that its correlation with bargaining power is sensitive

²⁶While our estimate of θ lies below standard production-function estimates, which are typically closer to one, this reflects the short-run nature of the object we identify. Standard estimates rely on firm-level variation in inputs and outputs and need not recover the same object as our parameter, which is a strictly short-run parameter.

²⁷The implied elasticity is computed as $d \ln MC_{ih} / d \ln q_{ijh} = \frac{1-\theta}{\theta} \cdot x_{ijh}$, using the average buyer share of 0.25 from Table 1.

TABLE 4: Estimated Model Primitives

Panel A: Calibrated Parameters				
$\hat{\nu}$	$\hat{\gamma}$		$\hat{\rho}$	
4	0.5		10	
Panel B: Estimated Parameters (GMM)				
	(1)	(2)	(3)	(4)
Rel. bargaining power: $\ln \widehat{\frac{\phi}{1-\phi}}$	1.563 (0.051)		1.168 (0.043)	
Returns to scale ($\hat{\theta}$)	0.455 (0.004)	0.505 (0.004)	0.417 (0.005)	0.508 (0.006)
Constant		4.327 (0.361)		2.099 (0.231)
Longevity		-0.216 (0.043)		0.476 (0.090)
Number of HS10 transactions		-0.342 (0.030)		-0.115 (0.024)
Multiple HS10 dummy		0.004 (0.038)		0.113 (0.039)
Lagged outside option		-0.285 (0.028)		-0.242 (0.032)
None	Yes	Yes	No	No
$FE_h + FE_t + FE_j$	No	No	Yes	Yes
Observations	3,120,000			
Panel C: Implied Bargaining Powers ($\hat{\phi}$)				
Mean	0.827 (0.007)	0.931 (0.081)	0.763 (0.008)	0.904 (0.079)
Median	— —	0.957 (0.081)	— —	0.927 (0.079)

Notes: This table presents model estimates using our main estimation sample of U.S. imports of intermediate inputs and capital goods, 2001–2016. Panel A reports calibrated parameters: demand elasticity (ν), cost elasticity to foreign input prices (γ), and the elasticity of substitution across foreign varieties (ρ). Panel B reports GMM estimates. Columns (1) and (3) impose a constant ϕ , while Columns (2) and (4) allow heterogeneous bargaining power via κ . Specifications differ in fixed effects. Controls include relationship longevity, transaction intensity, the relative outside option (lagged), and a multi-product indicator. Panel C reports the mean and median implied bargaining power. Robust standard errors; Panel C uses the delta method. Instruments include HS10-level counts of exporters and importers and lagged bilateral shares (excluding the focal pair). All estimations use CIF duty-exclusive unit values. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

to how unobserved heterogeneity is handled. The coefficient on transaction frequency is consistently negative and significant, while the multiple-product indicator is positive but precisely estimated only in the fixed-effects specification. By contrast, the coefficient on the relative outside option is stable and precisely estimated across all specifications: importers hold less bargaining power when their supplier has a stronger outside option. Specifically, an increase in the supplier’s past sales to other buyers relative to the importer’s purchases from other suppliers is consistently associated with lower bargaining power for the importer.

Robustness We assess the robustness of the structural estimates to alternative sample definitions and model calibrations. Table 5 reports the results.

TABLE 5: Robustness of Model Estimates

	Baseline	Only Con- sumption	All Products	$\rho = 5$	$\varrho = 0.5$
	(1)	(2)	(3)	(4)	(5)
Bargaining power (ϕ)	0.827 (0.007)	0.718 (0.005)	0.765 (0.004)	0.877 (0.004)	0.856 (0.007)
Returns to scale (θ)	0.455 (0.004)	0.558 (0.003)	0.508 (0.002)	0.465 (0.002)	0.470 (0.003)
Observations	3,120,000	3,024,000	6,143,000	3,120,000	3,120,000

Notes: This table reports robustness checks for the structural estimates. Columns vary either the sample definition or the calibrated parameters. “Only Consumption” restricts the sample to consumption goods. “All Products” includes all product categories. Columns $\rho = 5$ and $\varrho = 0.5$ modify the calibrated elasticities. All columns report GMM estimates without fixed effects. All estimations use CIF duty-exclusive unit values. Observation counts are rounded in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

Restricting the sample to consumption goods only (column (2)) yields a slightly lower estimate of ϕ and a slightly higher estimate of θ . Column (3) reports estimates for the pooled sample, including all products, which lie in between, as expected. The model’s core mechanisms, therefore, operate in consumption goods transactions as well. We nonetheless retain the baseline focus on capital and intermediate inputs, for which the framework is more naturally suited.²⁸

On the calibration side, we consider two alternative parameter values. First, we lower the elasticity of substitution across foreign varieties to $\rho = 5$ (instead of 10), at the lower end of

²⁸Adapting the framework specifically to consumption goods would likely require modifications to the modeling of the importer’s technology and its calibration.

the range reported in [Anderson and van Wincoop \(2004\)](#). Second, we introduce decreasing returns to scale in downstream production by setting $\varrho = 0.5$, close to the supplier-side estimate. Both ρ and ϱ affect how cross-sectional price dispersion varies with the exporter’s supplier share, and thus with oligopoly power. Columns (4) and (5) show that the structural estimates remain highly stable.²⁹ This suggests that ϕ is primarily identified from variation in buyer share rather than supplier share.

Implied Markups Using equation (2.10), we compute markups for all buyer-supplier–product matches given the estimated parameters and the observed market shares.

The resulting markup distribution is clustered near the competitive benchmark, with a median of 1.01. This contrasts with the oligopoly benchmark, whose median is 1.21 and whose interquartile range spans 1.11 to 2.45. In the model, observed markups are a convex combination of oligopoly and oligopsony markups, where the median oligopsony markup is 0.95. The large estimated bargaining power of buyers ($\hat{\phi} \approx 0.8$) shifts weight toward oligopsony, allowing importers to capture a substantial share of the surplus from exporters.

4.3 Model Validation

Proposition 2 yields a closed-form expression for the pass-through elasticity, which we use to assess the model’s out-of-sample performance. Specifically, we examine whether the model can predict the level and heterogeneity of price responses to the 2017–2018 tariff increases, which are excluded from estimation. Using observed shares in 2017 and the estimated parameters, the model-implied change in duty-exclusive prices is

$$\Delta \widehat{\ln p_{ijht}^*} = \Phi_{ijht}^* \left(s_{ijht}, x_{ijht} \mid \hat{\Theta} \right) \cdot \Delta \ln \left(1 + \tau_{c(i)ht} \right), \quad (4.5)$$

where τ_{cht} denotes the statutory tariff rate and Φ_{ijht}^* is the model-implied pass-through elasticity into duty-exclusive prices.³⁰ Because the model delivers a partial-equilibrium mapping from tariffs to prices while the data reflect equilibrium outcomes, this exercise constitutes a joint test of the model’s partial-equilibrium structure and its underlying mechanisms.

Table 6 reports pass-through elasticities in the data and in the model for the baseline specification (equation (3.2)) and three restricted variants, using baseline fixed effects.³¹ Col-

²⁹Results are robust to alternative related-party trade classifications. See Appendix D.3 and Table D.6 for details.

³⁰Proposition 2 characterizes pass-through into duty-inclusive prices; pass-through into duty-exclusive prices follows as $\Phi^* = \Phi - 1$.

³¹For the restricted variants, equation (4.5) is evaluated using the baseline parameter estimates from column (1) of Table 4 without re-estimation, except where parameters are restricted by construction (e.g.,

umn (1) reports the empirical estimates. Column (2) corresponds to the baseline model. Column (3) shuts down both bargaining and decreasing returns by imposing $\phi = 0$ and $\theta = 1$, yielding a standard Nash–Bertrand model with constant marginal costs, as in [Dhyne et al. \(2022\)](#). Column (4) sets $\theta = 1$, retaining bargaining but eliminating decreasing returns, as in [Gopinath and Itskhoki \(2011\)](#). Column (5) sets $\phi = 0$, eliminating importer bargaining power by assigning all pricing power to exporters, while preserving decreasing returns.

Panel A of Table 6 reports average pass-through of about 87% in the data. The baseline model predicts a lower pass-through, around 78%, yet remains close to the data with overlapping confidence intervals. The restricted variants also generate incomplete pass-through, to varying degrees: shutting down both bargaining and decreasing returns ($\phi = 0, \theta = 1$) yields about 91%; eliminating importer bargaining alone ($\phi = 0$) produces the lowest pass-through, around 72%; removing decreasing returns alone ($\theta = 1$) raises pass-through substantially. While all variants generate incomplete pass-through, confidence intervals do not reject any of them, limiting the informativeness of these differences for model fit.

As illustrated in Figure 2, pass-through reflects the balance of the cost and markup channels, and different (ϕ, θ) combinations can generate similar averages through different mechanisms, particularly in the relevant range of buyer and supplier shares. In the baseline, strategic substitutabilities partly offset the cost channel (Column (2)); setting $\phi = 0$ shifts the markup channel toward complementarities, further lowering pass-through and moving the model farther from the data (Column (5)); eliminating the cost channel alone still yields incomplete pass-through under strategic complementarities (Column (3)); and removing both the cost channel and the strategic complementarities produces near-complete pass-through (Column (4)). Average pass-through is therefore a weak discriminator across models, with more informative evidence coming from cross-sectional moments.

Panel B of Table 6 interacts tariff changes with lagged supplier and buyer shares. In the data, pass-through declines sharply with buyer share and increases modestly with supplier share. The baseline model closely replicates both gradients, with magnitudes comparable to those in the data. This alignment also sheds new light on Figure D.1: the approximately monotonic decline in pass-through across buyer-share quartiles, which we documented but did not interpret structurally in Section 3.4, finds a natural counterpart in the model, where the cost channel scales linearly with x_{ij} (Equation 2.16).

Eliminating decreasing returns ($\theta = 1$, Columns (3)–(4)) substantially attenuates the buyer-share gradient and counterfactually predicts lower pass-through for exporters with larger supplier shares. By contrast, setting $\phi = 0$ (Column (5)) yields cross-sectional patterns

$\phi = 0$ or $\theta = 1$).

TABLE 6: Price Responses and Relationship Heterogeneity: Data vs. Model

Panel A: Main Effect					
	Data	Baseline	$\phi = 0,$ $\theta = 1$	$\theta = 1$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln(1 + \tau_{cht})$	-0.126 (0.091)	-0.222 (0.008)	-0.094 (0.004)	-0.013 (0.001)	-0.279 (0.008)
R-squared	0.04	0.31	0.22	0.21	0.34
Panel B: Including Interactions					
	Data	Baseline	$\phi = 0,$ $\theta = 1$	$\theta = 1$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln(1 + \tau_{cht})$	-0.032 (0.091)	-0.145 (0.008)	-0.049 (0.009)	-0.005 (0.001)	-0.219 (0.009)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$	0.033 (0.080)	0.125 (0.013)	-0.079 (0.014)	-0.011 (0.003)	0.122 (0.011)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$	-0.407 (0.112)	-0.421 (0.023)	-0.046 (0.023)	-0.013 (0.003)	-0.366 (0.022)
R-squared	0.04	0.44	0.23	0.23	0.41
$FE_{ht} + FE_{cs}$	Yes	Yes	Yes	Yes	Yes
Observations	249,000				

Notes: This table reports the response of duty-exclusive CIF prices to tariff changes at the exporter–importer–product level. Column (1) presents reduced-form estimates using observed price changes in the data. Columns (2)–(5) report the corresponding responses implied by model-predicted price changes under each specification. Panel A reports the main effect of tariff changes, while Panel B allows the response to vary with lagged supplier shares (s_{ijht-1}) and buyer shares (x_{ijht-1}). All specifications include product–time and country–sector fixed effects ($FE_{ht} + FE_{cs}$). Standard errors are clustered at the HS8 product and exporter–country level. Observation counts are rounded to four significant digits per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

similar to the baseline. This is expected, given Figure 2: when decreasing returns are strong, pass-through heterogeneity is largely governed by the cost channel, leaving limited scope for bargaining power to shape the cross-sectional gradient.

Taken together, the baseline is the only specification consistent with both panels. Column (3) comes closest to the data on average (91%), but is strongly rejected on cross-sectional grounds. Column (5) matches the cross-sectional patterns but noticeably overstates price absorption relative to both the data and the baseline. The model’s ability to reproduce observed patterns of tariff pass-through lends support to its underlying mechanisms and suggests that the partial equilibrium framework captures key forces governing price adjustment in firm-to-firm trade.

Additional Validation. Table 6 uses baseline fixed effects throughout; Table D.7 in Appendix D.4 repeats the exercise under more stringent fixed effects and yields the same conclusions, both in terms of average pass-through and cross-sectional patterns.

One concern is that observed price changes may reflect shocks unrelated to tariffs, potentially leading the data to reject (or accept) a correctly (or incorrectly) specified model for spurious reasons. To address this, we formalize the comparison using the IV goodness-of-fit test of Adão et al. (2023), which isolates tariff-driven variation in observed prices. Figure D.2 in Appendix D.4 reports the results. Consistent with Panel A of Table 6, none of the variants can be formally rejected based on average pass-through alone. Among the non-rejected specifications, the baseline provides a closer fit than Column (5), consistent with Table 6. Table D.8 shows that these results are robust to alternative calibrations of the model parameters and remain stable across various fixed-effects specifications.

Finally, we assess the model’s ability to predict bilateral quantity and sales responses.³² Table D.9 in Appendix D.4 reports the results for quantities. Across specifications, tariff increases are associated with declines in traded volumes in both the data and the model, with larger effects under more demanding fixed effects. Interaction coefficients with buyer and supplier shares are imprecisely estimated in the data and sensitive to specification, suggesting limited power to detect heterogeneous effects. By contrast, the model predicts systematic patterns. Specifically, quantity responses become less negative at higher buyer shares, reflecting the smaller price increases faced by high-share buyers.

Appendix Table D.10 extends the IV tests to quantities and sales jointly. Consistent with

³²Predicted quantity and sales changes are given by $\Delta \ln \widehat{q_{ijht}} = -\widehat{\varepsilon}_{ijht} \Delta \ln \widehat{p_{ijht}}$ and $\Delta \ln \widehat{r_{ijht}} = -(1 - \widehat{\varepsilon}_{ijht}) \Delta \ln \widehat{p_{ijht}}$, where $\widehat{\varepsilon}_{ijht}$ is the match-specific residual demand elasticity computed from observed shares and estimated parameters.

our price results, the baseline and decreasing-returns specifications provide the best fit to the data, although these tests formally reject all model variants. This comparatively weaker fit for quantities does not necessarily undermine the model’s mechanism. Quantities are measured with significantly greater error in customs data, and the mapping from prices to quantities relies on estimated demand elasticities that may amplify this noise. Despite these measurement challenges, the results remain qualitatively consistent with the model’s core implication: prices govern the allocation of trade, with quantities adjusting along the demand curve.

5 Aggregate Pass-Through in Firm-to-Firm Trade

We use the estimated model to quantify the price impact of the 2018 tariffs within continuing buyer-supplier pairs and its underlying mechanisms. To aggregate the model’s micro-level predictions, we compute model-implied bilateral price changes and estimate a weighted pass-through for this subset of trade by regressing predicted price changes on the tariff shock, weighting observations by their initial import shares. We then use the model’s expression for cost and markup elasticities to decompose predicted price changes into its two underlying determinants. For each channel, we recompute counterfactual price changes while holding the other component fixed, and re-estimate the weighted regression.

Panel A of Table 7 reports aggregate pass-through predicted by the model, alongside a counterfactual that isolates the cost channel by shutting down markup adjustments. The model predicts aggregate pass-through around 73% under the baseline specification (Column (1)) and 67% under the specification with product–time, country–time, and buyer–time fixed effects (Column (2)), somewhat below the unweighted average of 78% reported in Column (2) of the top panel of Table 6, reflecting the greater weight placed on large, high-share buyers, which exhibit more incomplete pass-through.

Marginal cost adjustments account for most of the aggregate response. Shutting down markup elasticities ($\Gamma_{ij} = 0$) yields pass-through rates between 71.5% and 64.8%, very close to the aggregate estimates. This indicates that the cost channel accounts for most of the aggregate response. Markup adjustments operate in the opposite direction, exerting only a small offsetting effect and therefore raising pass-through only slightly relative to the cost-only benchmark.

Panel B provides a complementary decomposition by analyzing the variance of $\Lambda_{ij} + \Gamma_{ij}$, which governs pass-through through $1/(1 + \Lambda_{ij} + \Gamma_{ij})$. By construction, the contributions of cost and markup elasticities sum to one. Consistent with Panel A, nearly all cross-sectional

TABLE 7: Aggregate Tariff Pass-Through and Decomposition

	Baseline FE	Stringent FE
	(1)	(2)
<i>Panel A: Aggregate Pass-through (%)</i>		
Agg. pass-through: $1/(1 + \Lambda_{ij} + \Gamma_{ij})$	72.7	66.5
Cost channel only: $1/(1 + \Lambda_{ij})$	71.5	64.8
<i>Panel B: Variance Decomposition of $\Lambda_{ij} + \Gamma_{ij}$</i>		
Cost Elasticity: Λ_{ij}	1.023	0.997
Markup Elasticity: Γ_{ij}	-0.023	0.003

Notes: This table reports model-implied aggregate pass-through estimates following the 2018 U.S. tariff increases. Column (1) includes product-time and country-sector fixed effects. Column (2) includes product-time, country-time, and buyer-time fixed effects. The overall pass-through elasticity is computed as $1 +$ the estimated coefficient on $\Delta \ln(1 + \tau)$, and decomposed into contributions from the cost channel (Λ_{ij}) and the markup channel (Γ_{ij}). The counterfactual “Cost channel only” row shows the predicted pass-through when markup elasticities are set to zero. Panel B reports the relative contribution of each channel to the cross-sectional variance of $\Lambda_{ij} + \Gamma_{ij}$, the total elasticity governing pass-through. These shares sum to one and are derived from a variance decomposition. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

variation in aggregate pass-through is accounted for by the cost elasticity Λ_{ij} , with markup elasticities playing only a minor role.

5.1 Comparison with Existing Pass-Through Estimates

Several influential studies document near-complete pass-through of the 2018 U.S. tariffs at the product level (e.g., Fajgelbaum et al., 2020; Amiti et al., 2019b, 2020). Within our baseline sample of repeated firm-to-firm relationships, the implied average (unweighted) pass-through is bounded between 81 and 87 percent, substantially incomplete and closely matched by the model (see Appendix Tables 6 and D.7).³³

Two features distinguish our setting from product-level studies. First, our baseline imposes sample restrictions: we focus on repeated, arm’s-length links trading intermediate and capital goods, with suppliers serving at least two buyers. Second, we measure prices at the buyer-supplier-product level and at annual frequency, whereas product-level studies aggregate across relationships and use monthly data. To assess how each contributes to the gap with

³³In a specification comparable to those in the literature, Column (1) of Table 3 yields a duty-exclusive pass-through elasticity of -0.13 (s.e. 0.09); a more stringent set of fixed effects yields -0.19 (Appendix Table D.2).

the literature, we evaluate them in turn.

Sample restrictions. To isolate the role of sample composition, we progressively relax our baseline restrictions while keeping the unit of observation at the relationship level. Appendix Table D.11 shows the results. Including consumption goods raises pass-through from 82% to 88%; allowing single-buyer suppliers raises it further to about 92%; and the broadest specification, which also restores related-party trade, energy goods, and extreme price changes, reaches 95%. Each restriction contributes, but none on its own closes the gap with the literature: within the sample of repeated supplier–buyer links, pass-through remains incomplete regardless of which other restrictions are relaxed.

Aggregation. To isolate the role of aggregation, we aggregate our baseline firm-to-firm prices to the product–source-country level and use monthly rather than annual data, following the literature. Appendix D.5 provides more details on the data construction, and Table D.12 reports the results. Within our baseline sample of repeated links, pass-through remains incomplete even at the product level, confirming that aggregation alone does not close the gap. Similarly, pass-through remains incomplete even once all sample restrictions are lifted at the product level (Columns (1)–(4)). The remaining difference is accounted for by the final step: once we add non-consecutive relationships (Column (5)), that is, transactions not observed in both years, pass-through reaches completeness, matching the literature.

The near-complete pass-through in product-level studies is thus driven primarily by the inclusion of non-consecutive, one-off relationships. Within the population of ongoing buyer-supplier links, pass-through is systematically incomplete. Understanding aggregate pass-through therefore requires accounting for the extensive margin of trade, that is, how the set of active trading relationships responds to tariffs, beyond pricing within existing relationships.

6 Conclusions

We develop and estimate a structural model of two sided market power in firm-to-firm trade to shed light on tariff pass-through in global supply chains. The framework features oligopoly power on the exporter side and oligopsony power on the importer side, and yields closed-form expressions for bilateral markups and pass-through as functions of two sufficient statistics: the exporter’s supplier share and the importer’s buyer share, facilitating the mapping from theory to data, structural estimation, and the decomposition of pass-through mechanisms.

The two primitives governing oligopsony power, namely the inverse export supply elasticity and bargaining power, also govern the determinants of pass-through: the inverse export sup-

ply elasticity determines whether pass-through reflects cost or markup adjustments, while bargaining power shapes the nature of strategic markup adjustment. The framework therefore nests as special cases both the competitive benchmark under decreasing returns and the standard oligopoly pass-through model under constant returns.

Applied to U.S. imports in the 2018 tariff episode, our estimates reveal substantial oligopsony power and steep export supply curves. At these parameters, the cost channel dominates, explaining why pass-through within continuing buyer-supplier relationships is substantially incomplete and declines systematically with buyer share. This specific result likely reflects the relationship specificity and supply-side frictions that characterize U.S. firm-to-firm trade. In other settings, with weaker relationship specificity or flatter export supply, oligopsony power is lower and the framework accommodates a larger role for strategic markup adjustment.

Incomplete pass-through through the cost channel is thus a feature of short-run pricing in ongoing firm-to-firm relationships. Over longer horizons, as firms reallocate across partners and capacity constraints relax, effective returns to scale rise toward constant and the markup channel prevails, in line with the firm-level pass-through literature that attributes incomplete pass-through to strategic markup adjustment (e.g., [Amiti et al., 2014](#)). Even there, our framework provides a richer characterization of how bargaining power shapes pass-through across bilateral relationships. Sample composition matters too: product-level studies pool repeated arm's-length links with one-off and newly formed relationships, and we show that this extensive-margin variation accounts for the bulk of the near-complete pass-through documented at the product level ([Amiti et al., 2019b](#); [Fajgelbaum et al., 2020](#)).

Our analysis is explicitly partial equilibrium. While we show that this approach yields accurate predictions of pass-through within firm-to-firm relationships, a full account of tariff incidence requires accounting for extensive-margin adjustments in the set of active relationships, as well as general-equilibrium responses in wages and export prices. Whether the cost-channel dominance we document at the intensive margin survives once these forces are taken into account is a key open question for future research.

References

- ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): “The Network Origins of Aggregate Fluctuations,” *Econometrica*, 80, 1977–2016.
- ACEMOGLU, D. AND A. TAHBAZ-SALEHI (2025): “The Macroeconomics of Supply Chain Disruptions,” *Review of Economic Studies*, 92, 656–695.
- ADÃO, R., A. COSTINOT, AND D. DONALDSON (2023): “Putting quantitative models to the test: An application to trump’s trade war,” *NBER Working Paper #31321*.
- AMITI, M., O. ITSKHOKI, AND J. KONINGS (2014): “Importers, exporters, and exchange rate disconnect,” *The American Economic Review*, 104, 1942–1978.
- (2019a): “International shocks, variable markups, and domestic prices,” *The Review of Economic Studies*, 86, 2356–2402.
- AMITI, M., S. J. REDDING, AND D. E. WEINSTEIN (2019b): “The Impact of the 2018 Tariffs on Prices and Welfare,” *Journal of Economic Perspectives*, 33, 187–210.
- (2020): “Who’s Paying for the U.S. Tariffs? A Longer-Term Perspective,” *American Economic Review: Papers & Proceedings*, 110, 541–546.
- ANDERSON, J. E. AND E. VAN WINCOOP (2004): “Trade Costs,” *Journal of Economic Literature*, 42, 691–751.
- ANTRÀS, P. (2015): *Global production: Firms, contracts, and trade structure*, Princeton University Press.
- ATKESON, A. AND A. BURSTEIN (2008): “Trade Costs, Pricing to Market, and International Relative Prices,” *American Economic Review*, 98, 1998–2031.
- ATKIN, D., J. BLAUM, P. D. FAJGELBAUM, AND A. OSPITAL (2024): “Trade Barriers and Market Power: Evidence from Argentina’s Discretionary Import Restrictions,” *NBER Working Paper, #32037*.
- AUER, R. A. AND R. S. SCHOENLE (2016): “Market Structure and Exchange Rate Pass-Through,” *Journal of International Economics*, 98, 60–77.
- AVIGNON, R., C. CHAMBOLLE, E. GUIGUE, AND H. MOLINA (2024): “Markups, Mark-downs, and Bargaining in a Vertical Supply Chain,” *Available at SSRN*.
- AZKARATE-ASKASUA, M. AND M. ZERECERO (2025): “Union and firm labor market power,” *Available at SSRN 4323492*.
- BARROT, J. N. AND J. SAUVAGNAT (2016): “Input Specificity and the Propagation of Idiosyncratic Shocks in Production,” *Quarterly Journal of Economics*, 131, 1543 – 1592.
- BASU, S. AND J. G. FERNALD (1997): “Returns to scale in US production: Estimates and implications,” *Journal of political economy*, 105, 249–283.

- BERGER, D., K. HERKENHOFF, AND S. MONGEY (2022): “Labor market power,” *American Economic Review*, 112, 1147–1193.
- 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.
- BERNARD, A. B., J. B. JENSEN, AND P. K. SCHOTT (2006): “Transfer pricing by US-based multinational firms,” *NBER Working Paper*, #12493.
- BERNARD, A. B., A. MOXNES, AND Y. U. SAITO (2019): “Production networks, geography, and firm performance,” *Journal of Political Economy*, 127, 639–688.
- BOEHM, C. E., A. FLAAEN, AND N. PANDALAI-NAYAR (2019): “Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake,” *Review of Economics and Statistics*, 101, 60–75.
- BOEHM, C. E. AND N. PANDALAI-NAYAR (2022): “Convex supply curves,” *American Economic Review*, 112, 3941–3969.
- BRODA, C., N. LIMA, AND D. E. WEINSTEIN (2008): “Optimal tariffs and market power: the evidence,” *American Economic Review*, 98, 2032–65.
- BRODA, C. AND D. E. WEINSTEIN (2006): “Globalization and the Gains from Variety,” *The Quarterly Journal of Economics*, 121, 541–585.
- CARVALHO, V. M., M. NIREI, Y. U. SAITO, AND A. TAHBAZ-SALEHI (2021): “Supply chain disruptions: Evidence from the great east japan earthquake,” *The Quarterly Journal of Economics*, 136, 1255–1321.
- CAVALLO, A., G. GOPINATH, B. NEIMAN, AND J. TANG (2020): “Tariff passthrough at the border and at the store: evidence from US trade policy,” *American Economic Review: Insights*.
- COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. S. LEE (2019): ““Nash-in-Nash” Bargaining: A Microfoundation for Applied Work,” *Journal of Political Economy*, 127.
- DEMIRER, M. AND M. RUBENS (2025): “Welfare Effects of Buyer and Seller Power,” *NBER Working Paper* #33371.
- DHYNE, E., A. K. KIKKAWA, AND G. MAGERMAN (2022): “Imperfect Competition in Firm-to-Firm Trade,” *Journal of the European Economic Association*.
- EDMOND, C., V. MIDRIGAN, AND D. Y. XU (2023): “How costly are markups?” *Journal of Political Economy*, 131, 1619–1675.
- ELDRIDGE, L. AND S. POWERS (2018): “Imported Inputs to U.S. Production and Productivity: Two decades of Evidence,” *BLS Working Papers*.

- 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.
- FLAAEN, A. B., A. HORTAÇSU, AND F. TINTELNOT (2020): “The production relocation and price effects of US trade policy: the case of washing machines,” *American Economic Review*.
- FONTAINE, F., J. MARTIN, AND I. MEJEAN (2020): “Price Discrimination Within and Across EMU Markets: Evidence from French Exporters,” *Journal of International Economics*.
- GARETTO, S. (2016): “Firms’ heterogeneity, incomplete information, and pass-through,” *Journal of International Economics*, 101, 168–179.
- GOLDBERG, L. S. AND C. TILLE (2013): “A Bargaining Theory of Trade Invoicing and Pricing,” *NBER Working Paper*, #18985.
- GOLDBERG, P. K. AND R. HELLERSTEIN (2013): “A structural approach to identifying the sources of local currency price stability,” *The Review of Economic Studies*, 80, 175–210.
- GOPINATH, G. AND O. ITSKHOKI (2011): “In search of real rigidities,” *NBER Macroeconomics Annual*, 25, 261–310.
- GRASSI, B. (2018): “IO in I-O : Competition and Volatility in Input-Output Networks,” *mimeo*.
- HEISE, S. (2024): “Firm-to-firm relationships and the pass-through of shocks: Theory and evidence,” *Review of Economics and Statistics*, 1–45.
- HORN, H. AND A. WOLINSKY (1988): “Bilateral Monopolies and Incentives for Merger Author,” *RAND Journal of Economics*, 19, 408–419.
- HUANG, H., K. MANOVA, O. PERELLO, AND F. PISCH (2024): “Firm heterogeneity and imperfect competition in global production networks,” .
- IGNATENKO, A. (2024): “The Market Power of Buyers and Sellers in International Trade,” *mimeo*.
- KAMAL, F. AND R. MONARCH (2018): “Identifying Foreign Suppliers in U.S. Import Data,” *Review of International Economics*, 26, 117–139.
- MACEDONI, L. AND E. MATTANA (2024): “Pricing in firm-to-firm trade: evidence from a Danish multinational,” *Review of World Economics*, 160, 311–375.
- MORLACCO, M. (2019): “Market Power in Input Markets : Theory and Evidence from French Manufacturing,” *mimeo*.
- NUNN, N. (2007): “Relationship-Specificity, Incomplete Contracts, and the Pattern of Trade,” *Quarterly Journal of Economics*, 122, 569–600.
- PIERCE, J. R. AND P. K. SCHOTT (2009): “A Concordance between Ten-Digit U.S. Har-

- monized System Codes and Sic/Naics Product Classes and Industries,” *NBER Working Paper #15548*.
- RUHL, K. J. (2015): “How Well Is US Intrafirm Trade Measured?” *American Economic Review*, 105, 524–29.
- SANDERSON, E. AND F. WINDMEIJER (2016): “A weak instrument F-test in linear IV models with multiple endogenous variables,” *Journal of Econometrics*, 190, 212–221, endogeneity Problems in Econometrics.

ONLINE APPENDIX

“Two-Sided Market Power in Firm-to-Firm Trade”

Vanessa Alviarez, Michele Fioretti, Ken Kikkawa, Monica Morlacco

This Appendix is organized as follows. Section [A](#) contains derivations of all mathematical expressions in the text, including proofs of Propositions. Section [B](#) contains the discussion of theory extensions. Section [C](#) contains details about the data and sample selection criteria. Section [D](#) provides additional empirical results referenced in the text. Section [E](#) provides additional details regarding the estimation.

A Mathematical Derivations

A.1 Proof of Proposition 1

To derive μ_{ij} , we first obtain the demand elasticity ε_{ij} from the importer's cost minimization problem, then solve the bilateral bargaining problem.

Demand Elasticity ε_{ij}

Our baseline model assumes that the importer first chooses the input quantity to minimize costs given a price, then negotiates the price bilaterally with the exporter. The importer's nested CES production structure (equations (2.2)–(2.3)) yields the demand for input q_{ij} in (2.6):

$$q_{ij} = q_j^f \varsigma_{ij}^\rho \left(\frac{p_{ij}}{p_j^f} \right)^{-\rho},$$

where $p_j^f = (\sum_i \varsigma_{ij}^\rho p_{ij}^{1-\rho})^{\frac{1}{1-\rho}}$ is the price index (i.e., shadow cost) of imported inputs. Solving the outer CES aggregator for total input demand yields:

$$q_j^f = \gamma MC_j q_j (p_j^f)^{-1}, \quad (\text{A.1})$$

$$q_j^d = (\varrho - \gamma) MC_j q_j (p_j^d)^{-1}, \quad (\text{A.2})$$

where MC_j is the unit cost of output q_j , given by:

$$MC_j = \left[\varphi_j^{-1} \left(\frac{p_j^f}{\gamma} \right)^\gamma \left(\frac{p_j^d}{\varrho - \gamma} \right)^{\varrho - \gamma} \right]^{\frac{1}{\varrho}} q_j^{\frac{1-\varrho}{\varrho}} \equiv k_j q_j^{\frac{1-\varrho}{\varrho}}. \quad (\text{A.3})$$

This form illustrates symmetry with exporter-side technology. Finally, note that the cost share of foreign inputs is constant and equal to

$$\frac{p_j^f q_j^f}{p_j^f q_j^f + p_j^d q_j^d} = \frac{\gamma}{\varrho}.$$

A relevant object for our derivations will be $\frac{d \ln MC_j}{d \ln p_j^f}$, namely, the elasticity of the marginal cost MC_j with respect to p_j^f . To find this elasticity, we first use the demand function downstream to write q_j as

$$q_j = \left(\frac{\nu}{\nu - 1} \right)^{-\nu} MC_j^{-\nu} D_j, \quad (\text{A.4})$$

where D_j is the firm-level demand shifter. Substituting equation (A.4) into equation (A.3) and rearranging, we can write:

$$MC_j = \left[z_j^{-1} \left(\frac{p_j^f}{\gamma} \right)^\gamma \left(\frac{p_j^d}{\varrho - \gamma} \right)^{(\varrho - \gamma)} \right]^{\frac{1}{\varrho + \nu - \nu\varrho}} \left(\frac{\nu}{\nu - 1} \right)^{-\nu \frac{1 - \varrho}{\varrho + \nu - \nu\varrho}} (D_j)^{\frac{1 - \varrho}{\varrho + \nu - \nu\varrho}},$$

which implies

$$\frac{d \ln MC_j}{d \ln p_j^f} = \frac{\gamma}{\varrho + \nu - \nu\varrho}.$$

Armed with these equations, we proceed to find the elasticity of interest. Given the log demand:

$$\ln q_{ij} = \ln q_j^f + \rho \ln s_{ij} - \rho (\ln p_{ij} - \ln p_j^f),$$

and equation (A.1), we find:

$$\begin{aligned} \varepsilon_{ij} &= -\frac{d \ln q_{ij}}{d \ln p_{ij}} = -\left(\frac{d \ln q_j^f}{d \ln p_j^f} + \rho \right) \frac{d \ln p_j^f}{d \ln p_{ij}} + \rho \\ &= \left((\nu - 1) \frac{d \ln c_j}{d \ln p_j^f} + 1 - \rho \right) s_{ij} + \rho \\ &= \eta s_{ij} + (1 - s_{ij})\rho, \end{aligned} \tag{A.5}$$

where we defined

$$\begin{aligned} \eta &\equiv -\frac{d \ln q_j^f}{d \ln p_j^f} = \frac{(\nu - 1)\gamma}{\varrho + \nu(1 - \varrho)} + 1. \\ &= \frac{(\varrho - \gamma) + \nu(1 - (\varrho - \gamma))}{\varrho + \nu(1 - \varrho)}. \end{aligned}$$

Equilibrium Price

The problem of the $i - j$ pair is to choose a bilateral price p_{ij} that solves the following problem:

$$\max_p \left(\underbrace{\pi_i(p) - \tilde{\pi}_{i(-j)}}_{GFT_{ij}^i(p)} \right)^{1-\phi} \left(\underbrace{\pi_j(p) - \tilde{\pi}_{j(-i)}}_{GFT_{ij}^j(p)} \right)^\phi, \tag{A.6}$$

where $\phi \in (0, 1)$ is j 's bargaining power, and the terms inside parentheses are the gains from trade for exporter i ($GFT_{ij}^i(p)$) and importer j ($GFT_{ij}^j(p)$), written as a function of p_{ij} .

The FOC associated with the problem (A.6) can be written as:

$$0 = \frac{d \ln \pi_i}{d \ln p_{ij}} + \frac{\phi}{1 - \phi} \cdot \frac{GFT_{ij}^i}{\pi_i} \frac{\pi_j}{GFT_{ij}^j} \cdot \frac{d \ln \pi_j}{d \ln p_{ij}}, \quad (\text{A.7})$$

where we used the fact that $\frac{dGFT_{ij}^k}{dp} = \frac{d\pi_k}{dp}$ for $k = \{i, j\}$. In what follows, we derive expressions for $\frac{d \ln \pi_i}{d \ln p_{ij}}$, $\frac{d \ln \pi_j}{d \ln p_{ij}}$, GFT_{ij}^i and GFT_{ij}^j .

Exporter i 's Profits and Gains from Trade—Firm i 's profit under a successful negotiation can be expressed as

$$\pi_i = p_{ij}q_{ij} + \sum_{k \neq j} p_{ik}q_{ik} - \theta MC_i q_i,$$

where $MC_i = k_i q_i^{\frac{1-\theta}{\theta}}$. The elasticity of the profit π_i with respect to p_{ij} can be found as:

$$\frac{d \ln \pi_i}{d \ln p_{ij}} = \frac{q_{ij}}{\pi_i} (\varepsilon_{ij} - 1) \left(\mu_{ij}^{\text{oligopoly}} MC_i - p_{ij} \right), \quad (\text{A.8})$$

where $\varepsilon_{ij} \equiv -\frac{d \ln q_{ij}}{d \ln p_{ij}}$ is defined in equation (A.5) and

$$\mu^{\text{oligopoly}} := \frac{\varepsilon_{ij}}{\varepsilon_{ij} - 1}.$$

The “outside” profit of firm i are

$$\tilde{\pi}_{i(-j)} = \sum_{k \neq j} p_{ik}q_{ik} - \theta \widetilde{MC_i q_i}.$$

For a constant k_i and given $q_i = \sum_{j \in \mathcal{Z}_i} q_{ij}$, the total cost in case of failed agreement can be found as:

$$\begin{aligned} \theta \widetilde{MC_i q_i} &= \theta MC_i q_i \left(\frac{\tilde{q}_i}{q_i} \right)^{\frac{1}{\theta}}. \\ &= \theta MC_i q_i (1 - x_{ij})^{\frac{1}{\theta}}. \end{aligned}$$

The exporter's gains from trade GFT_{ij}^i are thus given by:

$$\begin{aligned} GFT_{ij}^i(p_{ij}) &\equiv \pi_i(p) - \tilde{\pi}_{i(-j)} \\ &= q_{ij} \left(p_{ij} - MC_i \mu_{ij}^{\text{oligopsony}} \right), \end{aligned} \quad (\text{A.9})$$

where $\mu_{ij}^{\text{oligopsony}} := \theta \left[\frac{\Delta_{ij}^x}{x_{ij}} \right]$ and $\Delta_{ij}^x := \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]$.

Importer j 's Profits and Gains from Trade—Firm j 's profit under a successful negotiation can be expressed as

$$\pi_j = p_j q_j - \varrho MC_j q_j = \frac{\gamma}{\eta - 1} MC_j q_j,$$

where MC_j is defined in equation (A.3), and q_j in equation (A.4). The elasticity of the profit π_j with respect to p_{ij} can be found as:

$$-\frac{d \ln \pi_j}{d \ln p_{ij}} = (\eta - 1) s_{ij}. \quad (\text{A.10})$$

The outside profit of firm j under a failed negotiation is

$$\tilde{\pi}_{j(-i)} = \frac{\gamma}{\eta - 1} \widetilde{MC_j q_j},$$

where $\varrho \widetilde{MC_j q_j}$ denotes the total “outside” cost of importer j . The assumptions on technology and demand downstream imply that we can write:

$$\widetilde{MC_j q_j} = (1 - s_{ij})^{\frac{\eta-1}{\rho-1}} MC_j q_j.$$

Putting things together, the importer's gains from trade can be written as:

$$GFT_j(p_{ij}) = \pi_j \Delta_{ij}^s \quad (\text{A.11})$$

where we defined $\Delta_{ij}^s := \left(1 - (1 - s_{ij})^{\frac{\eta-1}{\rho-1}} \right)$.

We're now ready to solve for the bilateral price p_{ij} . Substituting equations (A.8), (A.9), (A.10), and (A.11) into the FOC in equation (A.7), we obtain:

$$\mu_{ij} := \frac{p_{ij}}{MC_i} = (1 - \omega_{ij}) \mu_{ij}^{\text{oligopoly}} + \omega_{ij} \mu_{ij}^{\text{oligopsony}} \quad (\text{A.12})$$

where the weighting factor is

$$\omega_{ij} := \frac{\frac{\phi}{1-\phi} \lambda_{ij}}{1 + \frac{\phi}{1-\phi} \lambda_{ij}} \in (0, 1) \quad (\text{A.13})$$

and where

$$\begin{aligned}
\lambda_{ij} &:= \frac{\pi_j}{GFT_{ij}^j} \cdot \left(-\frac{d \ln \pi_j}{d \ln p_{ij}} \cdot (\varepsilon_{ij} - 1)^{-1} \right) \\
&= \underbrace{\frac{\pi_j}{GFT_{ij}^j}}_{\lambda_{ij}^N} \cdot \underbrace{-\frac{d \ln \pi_j}{d \ln p_{ij} q_{ij}}}_{\lambda_{ij}^I} \\
&= (\Delta_{ij}^s)^{-1} \cdot \frac{\eta - 1}{\varepsilon_{ij} - 1} s_{ij} > 0.
\end{aligned} \tag{A.14}$$

Equations (A.12)-(A.14) forms the basis for Proposition 1 in Section 2. □

A.2 Proof of Proposition 2

Proposition 2 characterizes how the bilateral markup μ_{ij} varies with bilateral market shares. Part I states that a positive relationship between μ_{ij} and the exporter's supplier share s_{ij} requires $\phi < 1$. Part II states that a negative relationship between μ_{ij} and the importer's buyer share x_{ij} can arise only if $\phi > 0$ and $\theta < 1$.

Proof of Part I

The Nash bargaining solution implies that the bilateral markup satisfies

$$\mu_{ij} = (1 - \omega_{ij})\mu^{\text{oligopoly}}(s_{ij}) + \omega_{ij}\mu^{\text{oligopsony}}(x_{ij}), \quad (\text{A.15})$$

with terms defined above. Differentiating with respect to s_{ij} yields

$$\frac{\partial \mu_{ij}}{\partial s_{ij}} = (1 - \omega_{ij}) \frac{\partial \mu_{ij}^{\text{oligopoly}}}{\partial s_{ij}} + \frac{\partial \omega_{ij}}{\partial s_{ij}} (\mu_{ij}^{\text{oligopsony}} - \mu_{ij}^{\text{oligopoly}}). \quad (\text{A.16})$$

Under the maintained assumption $\eta < \rho$, the oligopoly markup is strictly increasing in supplier share, so $\partial_s \mu_{ij}^{\text{oligopoly}} > 0$, while the oligopsony markdown is strictly lower than the oligopoly markup, $\mu_{ij}^{\text{oligopsony}} - \mu_{ij}^{\text{oligopoly}} < 0$.

Consider first the case $\phi = 1$, which implies $\omega_{ij} = 1$ for all s_{ij} . Hence, $(1 - \omega_{ij}) = 0$ and $\partial \omega_{ij} / \partial s_{ij} = 0$. Both terms in (A.16) therefore vanish, so that

$$\frac{\partial \mu_{ij}}{\partial s_{ij}} = 0 \quad \text{when } \phi = 1.$$

Thus, a strictly positive relationship between μ_{ij} and s_{ij} cannot arise under full importer bargaining power.

Now suppose $\phi < 1$. In this case $\omega_{ij}(s_{ij})$ is hump-shaped in s_{ij} with $\omega_{ij}(0) = \omega_{ij}(1) = \phi$, implying that there exists $\bar{s} \in (0, 1)$ such that $\partial \omega_{ij} / \partial s_{ij} < 0$ for all $s_{ij} > \bar{s}$. For such values of s_{ij} , the second term in (A.16) is nonnegative, since it is the product of two negative terms. Together with $(1 - \omega_{ij}) > 0$ and $\partial_s \mu_{ij}^{\text{oligopoly}} > 0$, this implies

$$\frac{\partial \mu_{ij}}{\partial s_{ij}} > 0 \quad \text{for all } s_{ij} > \bar{s}.$$

For $s_{ij} < \bar{s}$, the sign of $\partial \mu_{ij} / \partial s_{ij}$ depends on the relative strength of the two terms in (A.16). In particular, $\partial \mu_{ij} / \partial s_{ij} > 0$ if and only if the direct oligopoly effect dominates the

bargaining-weight effect, i.e.,

$$(1 - \omega_{ij}) \frac{\partial \mu_{ij}^{\text{oligopoly}}}{\partial s_{ij}} > \left| \frac{\partial \omega_{ij}}{\partial s_{ij}} (\mu_{ij}^{\text{oligopsony}} - \mu_{ij}^{\text{oligopoly}}) \right| > 0.$$

We conclude that $\partial \mu_{ij} / \partial s_{ij} > 0$ can arise only when $\phi < 1$. This establishes Part I of Proposition 2. $\square$

Proof of Part II.

From (A.15), the bilateral markup μ_{ij} depends on the importer's buyer share x_{ij} only through the oligopsony markdown $\mu_{ij}^{\text{oligopsony}}(x_{ij})$. Differentiating with respect to x_{ij} :

$$\frac{\partial \mu_{ij}}{\partial x_{ij}} = \omega_{ij} \frac{\partial \mu_{ij}^{\text{oligopsony}}}{\partial x_{ij}}.$$

If $\phi = 0$, then $\omega_{ij} = 0$ for all x_{ij} , implying $\partial \mu_{ij} / \partial x_{ij} = 0$. Hence, a negative relationship between μ_{ij} and x_{ij} cannot occur if $\phi = 0$.

For $\phi > 0$, the sign of $\partial_x \mu_{ij}$ is equal to the sign of $\partial_x \mu_{ij}^{\text{oligopsony}}$. This is given by:

$$\frac{\partial \mu_{ij}^{\text{oligopsony}}}{\partial x_{ij}} = \frac{x_{ij}(1 - x_{ij})^{\frac{1}{\theta}-1} - \theta[1 - (1 - x_{ij})^{\frac{1}{\theta}}]}{x_{ij}^2}.$$

We establish the following lemma.

Lemma 1. *For any $x \in (0, 1)$ and $\theta > 0$,*

$$\text{sign}\left(x(1 - x)^{\frac{1}{\theta}-1} - \theta[1 - (1 - x)^{\frac{1}{\theta}}]\right) = \text{sign}(\theta - 1).$$

Proof. Let $t = (1 - x)^{1/\theta} \in (0, 1)$. Then the expression in the lemma can be written as

$$t^{1-\theta} - t - \theta + t\theta.$$

Express the above as a function of θ , $f(\theta)$. One can show that $f(\theta)$ is strictly convex in θ ($\partial^2 f(\theta) / \partial \theta^2 > 0$), and $f(0) = f(1) = 0$. This implies

$$t^{1-\theta} - t - \theta + t\theta \begin{cases} \leq 0 & \text{if } 0 \leq \theta \leq 1, \\ \geq 0 & \text{if } \theta \geq 1, \end{cases}$$

with equality at $\theta = 1$. □

By the lemma, $\partial \mu_{ij}^{\text{oligopsony}} / \partial x_{ij} < 0$ if and only if $\theta < 1$ since $x_{ij}^2 > 0$ for all $x \in (0, 1]$. Since $\omega_{ij} > 0$ when $\phi > 0$, this implies

$$\frac{\partial \mu_{ij}}{\partial x_{ij}} < 0 \iff \phi > 0 \text{ and } \theta < 1.$$

Therefore, observing a negative relationship between the bilateral markup and the importer's buyer share implies positive importer bargaining power and decreasing returns to scale, establishing Part II of Proposition 2. □

A.3 Proof of Proposition 3

The log (tariff-inclusive) price is given by:

$$\ln p_{ij} = \ln \mu_{ij} + \ln MC_i + \ln T_c,$$

where MC_i and μ_{ij} are as in equations (2.1) and (2.10), respectively.

Taking a full log-differential and rearranging terms yields:

$$\begin{aligned} d \ln p_{ij} &= -\Gamma_{ij} \cdot d \ln p_{ij} - \Lambda_{ij} d \ln p_{ij} + d \ln T_c \\ \frac{d \ln p_{ij}}{d \ln T_c} &= \frac{1}{1 + \Gamma_{ij} + \Lambda_{ij}}, \end{aligned}$$

where $\Gamma_{ij} \equiv -\frac{d \ln \mu_{ij}}{d \ln p_{ij}}$ and $\Lambda_{ij} \equiv -\frac{d \ln MC_i}{d \ln p_{ij}}$ are the partial *markup* and *cost* elasticities, respectively.

The Cost Elasticity

Taking the logarithm of equation (2.1), we obtain:

$$\ln MC_i = \ln k_i + \frac{1 - \theta}{\theta} \ln q_i.$$

It immediately follows that:

$$\begin{aligned} \Lambda_{ij} &\equiv -\frac{d \ln MC_i}{d \ln p_{ij}} = \frac{1 - \theta}{\theta} \frac{d \ln q_i}{d \ln q_{ij}} \left(-\frac{d \ln q_{ij}}{d \ln p_{ij}} \right) \\ &= \frac{1 - \theta}{\theta} \cdot x_{ij} \cdot \varepsilon_{ij} \geq 0. \end{aligned}$$

Moreover, the comparative statics with respect to the bilateral shares are easy to compute as:

$$\frac{d \Lambda_{ij}}{d x_{ij}} = \frac{1 - \theta}{\theta} \cdot \varepsilon_{ij} \geq 0,$$

with strict inequality whenever $\theta < 1$, whereas:

$$\frac{d \Lambda_{ij}}{d s_{ij}} = \frac{1 - \theta}{\theta} \cdot x_{ij} \cdot (\eta - \rho) < 0.$$

Thus, the cost elasticity weakly increases with the importer's buyer share x_{ij} , and it decreases with the exporter's supplier share s_{ij} .

Markup Elasticity

Taking logs of equation (2.10) and differentiating, we obtain:

$$\begin{aligned} d \ln \mu_{ij} = & \frac{(1 - \omega_{ij})\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}} d \ln \mu_{ij}^{\text{oligopoly}} + \frac{\omega_{ij}\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} d \ln \mu_{ij}^{\text{oligopsony}} \\ & + \frac{\omega_{ij} \left(\mu_{ij}^{\text{oligopsony}} - \mu_{ij}^{\text{oligopoly}} \right)}{\mu_{ij}} d \ln \omega_{ij}. \end{aligned}$$

Rearranging terms, the price elasticity of the bilateral markup can be expressed as:

$$\Gamma_{ij} \equiv -\frac{d \ln \mu_{ij}}{d \ln p_{ij}} = \frac{(1 - \omega_{ij})\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}} \Gamma_{ij}^{\text{oligopoly}} + \frac{\omega_{ij}\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} \Gamma_{ij}^{\text{oligopsony}} + \left(1 - \frac{\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}} \right) \Gamma_{ij}^{\omega},$$

where: $\Gamma_{ij}^{\text{oligopoly}} \equiv -\frac{d \ln \mu_{ij}^{\text{oligopoly}}}{d \ln p_{ij}}$, $\Gamma_{ij}^{\text{oligopsony}} \equiv -\frac{d \ln \mu_{ij}^{\text{oligopsony}}}{d \ln p_{ij}}$, and $\Gamma_{ij}^{\omega} \equiv -\frac{d \ln \omega_{ij}}{d \ln p_{ij}}$.

Oligopoly Markup Elasticity—The oligopoly markup elasticity is given by:

$$\Gamma_{ij}^{\text{oligopoly}} \equiv -\frac{d \ln \mu_{ij}^{\text{oligopoly}}}{d \ln p_{ij}} = -\frac{d \ln \mu_{ij}^{\text{oligopoly}}}{d \ln s_{ij}} \cdot \frac{d \ln s_{ij}}{d \ln p_{ij}}.$$

From the definition of s_{ij} , the last term is:

$$\frac{d \ln s_{ij}}{d \ln p_{ij}} = -(\rho - 1)(1 - s_{ij}).$$

Given $\mu_{ij}^{\text{oligopoly}} = \frac{\varepsilon_{ij}}{\varepsilon_{ij} - 1}$, we find:

$$-\frac{d \ln \mu_{ij}^{\text{oligopoly}}}{d \ln s_{ij}} = -\frac{1}{(\varepsilon_{ij} - 1)} \cdot \frac{\rho - \varepsilon_{ij}}{\varepsilon_{ij}},$$

which implies:

$$\Gamma_{ij}^{\text{oligopoly}} = -\frac{d \ln \mu_{ij}^{\text{oligopoly}}}{d \ln s_{ij}} \cdot \frac{d \ln s_{ij}}{d \ln p_{ij}} = \frac{1}{\varepsilon_{ij} - 1} \frac{\rho - \varepsilon_{ij}}{\varepsilon_{ij}} (\rho - 1)(1 - s_{ij}) \geq 0.$$

Oligopsony Markup Elasticity– The oligopsony markup elasticity is given by:

$$\Gamma_{ij}^{\text{oligopsony}} \equiv - \frac{d \ln \mu_{ij}^{\text{oligopsony}}}{d \ln p_{ij}} = - \left(\frac{d \ln \mu_{ij}^{\text{oligopsony}}}{d \ln x_{ij}} \right) \left(\frac{d \ln x_{ij}}{d \ln p_{ij}} \right).$$

From the definition of x_{ij} , the last term is:

$$\frac{d \ln x_{ij}}{d \ln p_{ij}} = -(1 - x_{ij})\varepsilon_{ij}.$$

Given $\mu_{ij}^{\text{oligopsony}} := \theta \left(\frac{1 - (1 - x_{ij})^{\frac{1}{\theta}}}{x_{ij}} \right)$, we find:

$$\frac{d \ln \mu_{ij}^{\text{oligopsony}}}{d \ln x_{ij}} = \left(\frac{x_{ij} (1 - x_{ij})^{\frac{1}{\theta} - 1}}{\theta \left(1 - (1 - x_{ij})^{\frac{1}{\theta}} \right)} - 1 \right)$$

which implies:

$$\Gamma_{ij}^{\text{oligopsony}} = - \left(\frac{d \ln \mu_{ij}^{\text{oligopsony}}}{d \ln x_{ij}} \right) \left(\frac{d \ln x_{ij}}{d \ln p_{ij}} \right) = \left(\frac{x_{ij} (1 - x_{ij})^{\frac{1}{\theta} - 1}}{\theta \left(1 - (1 - x_{ij})^{\frac{1}{\theta}} \right)} - 1 \right) (1 - x_{ij})\varepsilon_{ij}.$$

with

$$\begin{aligned} \frac{\partial \Gamma_{ij}^{\text{oligopsony}}}{\partial x_{ij}} &= \left[\frac{\partial}{\partial x_{ij}} \left(\frac{x_{ij} (1 - x_{ij})^{\frac{1}{\theta} - 1}}{\theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]} \right) (1 - x_{ij}) \right. \\ &\quad \left. - \left(\frac{x_{ij} (1 - x_{ij})^{\frac{1}{\theta} - 1}}{\theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]} - 1 \right) \right] \varepsilon_{ij} \\ &= \left[\frac{\left((1 - x_{ij})^{\frac{1-\theta}{\theta}} - \frac{1-\theta}{\theta} x_{ij} (1 - x_{ij})^{\frac{1}{\theta} - 2} \right) \theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right] - x_{ij} (1 - x_{ij})^{\frac{2}{\theta} - 2}}{\theta^2 \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]^2} \right] (1 - x_{ij}) \\ &\quad - \left(\frac{x_{ij} (1 - x_{ij})^{\frac{1}{\theta} - 1}}{\theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]} - 1 \right) \\ &= \left[\left(\frac{1 - \frac{1-\theta}{\theta} \cdot \frac{x_{ij}}{1 - x_{ij}}}{\theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]} - \frac{(1 - x_{ij})^{\frac{1}{\theta} - 1}}{\mu_{ij}^{\text{oligopsony}} \theta \left[1 - (1 - x_{ij})^{\frac{1}{\theta}} \right]} \right) (1 - x_{ij})^{\frac{1}{\theta}} \right] \end{aligned}$$

$$- \left(\frac{x_{ij}(1-x_{ij})^{\frac{1}{\theta}-1}}{\theta \left[1 - (1-x_{ij})^{\frac{1}{\theta}} \right]} - 1 \right) \Bigg]$$

Omega Elasticity

The elasticity of the weight ω_{ij} with respect to price is:

$$\Gamma_{ij}^{\omega} \equiv - \frac{d \ln \omega_{ij}}{d \ln p_{ij}} = \frac{d \ln \omega_{ij}}{d \ln s_{ij}} \left(- \frac{d \ln s_{ij}}{d \ln p_{ij}} \right).$$

Given

$$\frac{d \ln \omega_{ij}}{d \ln s_{ij}} = (1 - \omega_{ij}) \frac{d \ln \lambda_{ij}}{d \ln s_{ij}},$$

the above elasticity becomes:

$$\Gamma_{ij}^{\omega} = \frac{d \ln \lambda_{ij}}{d \ln s_{ij}} (1 - \omega_{ij}) (\rho - 1) (1 - s_{ij}),$$

Since $\lambda_{ij} = \frac{(\eta-1)s_{ij}}{\varepsilon_{ij}-1} \cdot \left(1 - (1-s_{ij})^{\frac{\eta-1}{\rho-1}} \right)^{-1}$, we find:

$$\frac{d \ln \lambda_{ij}}{d \ln s_{ij}} = 1 - \frac{\varepsilon_{ij} - \rho}{\varepsilon_{ij} - 1} - \frac{\eta - 1}{\rho - 1} \cdot \frac{(1 - s_{ij})^{\frac{\eta-1}{\rho-1}}}{1 - (1 - s_{ij})^{\frac{\eta-1}{\rho-1}}} \cdot \frac{s_{ij}}{(1 - s_{ij})}.$$

Thus:

$$\Gamma_{ij}^{\omega} = \left(1 - \frac{\varepsilon_{ij} - \rho}{\varepsilon_{ij} - 1} - \frac{\eta - 1}{\rho - 1} \cdot \frac{(1 - s_{ij})^{\frac{\eta-1}{\rho-1}}}{1 - (1 - s_{ij})^{\frac{\eta-1}{\rho-1}}} \cdot \frac{s_{ij}}{(1 - s_{ij})} \right) (1 - \omega_{ij}) (\rho - 1) (1 - s_{ij}).$$

□

A.4 Proof of Proposition 4

Proposition 4 states that if the tariff pass-through elasticity Φ_{ij} decreases with the importer's buyer share x_{ij} , then the exporter's technology exhibits decreasing returns to scale, i.e. $\theta < 1$.

Proof. The bilateral pass-through elasticity is

$$\Phi_{ij} = \frac{1}{1 + \Gamma_{ij} + \Lambda_{ij}},$$

where Γ_{ij} and Λ_{ij} denote the markup and marginal cost elasticities, respectively. Since $\partial\Phi_{ij}/\partial x_{ij}$ has the opposite sign of $\partial(\Gamma_{ij} + \Lambda_{ij})/\partial x_{ij}$, it suffices to show that

$$\partial(\Gamma_{ij} + \Lambda_{ij})/\partial x_{ij} > 0 \quad \text{for } \theta < 1$$

to prove the claim.

We first consider the case $\theta = 1$.

Under constant returns to scale, the exporter's marginal cost is constant and $\Lambda_{ij} = 0$. Moreover, the oligopsony markdown equals one, implying $\Gamma_{ij}^{\text{oligopsony}} = 0$. As Γ_{ij} depends on x_{ij} only through $\Gamma_{ij}^{\text{oligopsony}}$, it follows that $\partial\Gamma_{ij}/\partial x_{ij} = 0$ and hence $\partial\Phi_{ij}/\partial x_{ij} = 0$.

We now consider $\theta < 1$.

The cost elasticity is

$$\Lambda_{ij} = \frac{1 - \theta}{\theta} x_{ij} \varepsilon_{ij},$$

which is strictly increasing in x_{ij} and independent of the bargaining parameter ϕ . The markup elasticity depends on x_{ij} only through the oligopsony component $\Gamma_{ij}^{\text{oligopsony}}$. Varying ϕ therefore only scales the contribution of this term. Consequently, if

$$\left. \frac{\partial}{\partial x_{ij}} (\Gamma_{ij} + \Lambda_{ij}) \right|_{\phi=1} > 0,$$

where the contribution of x is maximized, then the same must hold for all $\phi \in [0, 1]$. It is therefore sufficient to study the case $\phi = 1$.

When $\phi = 1$, the oligopsony markdown elasticity can be written as

$$\Gamma_{ij}^{\text{oligopsony}} = f(x_{ij}) \varepsilon_{ij}, \quad f(x) = \left(\frac{x(1-x)^{\frac{1}{\theta}-1}}{\theta \left[1 - (1-x)^{\frac{1}{\theta}} \right]} - 1 \right) (1-x).$$

Hence,

$$\frac{\partial}{\partial x_{ij}} (\Gamma_{ij}^{\text{oligopsony}} + \Lambda_{ij}) = \varepsilon_{ij} \left[f'(x_{ij}) + \frac{1-\theta}{\theta} \right].$$

The function f is smooth on $(0, 1)$ and U-shaped: $f'(x) < 0$ for small x , $f'(x) > 0$ for large x , and $f'(x)$ is strictly increasing in x . Thus the minimum of $f'(x)$ over $(0, 1)$ is attained at the boundary as $x \rightarrow 0$.

A first-order expansion around $x = 0$ yields

$$f(x) = \frac{\theta-1}{2\theta}x + O(x^2), \quad \lim_{x \rightarrow 0} f'(x) = \frac{\theta-1}{2\theta}.$$

When $\theta < 1$, this implies

$$f'(x) \geq -\frac{1-\theta}{2\theta} > -\frac{1-\theta}{\theta} \quad \forall x \in (0, 1).$$

Therefore,

$$f'(x_{ij}) + \frac{1-\theta}{\theta} > 0 \quad \forall x_{ij} \in (0, 1),$$

and hence

$$\frac{\partial}{\partial x_{ij}} (\Gamma_{ij} + \Lambda_{ij}) > 0 \quad \forall \phi \in [0, 1] \quad \& \quad \theta < 1.$$

□

B Theory: Extensions

B.1 Efficient Bargaining

In the efficient bargaining setup, the importer and exporter jointly choose the price p_{ij} and quantity q_{ij} by maximizing the Nash product:

$$\max_{p_{ij}, q_{ij}} (GFT_{ij}^i)^{1-\phi} (GFT_{ij}^j)^\phi \quad \text{s.t.} \quad GFT_{ij}^i \geq 0, \quad GFT_{ij}^j \geq 0,$$

where $GFT_{ij}^k \equiv \pi_{ij}^k - \tilde{\pi}_{(-,.)}^k$, $k \in \{i, j\}$, are as defined in Section 2. This corresponds to a setting in which the firm pair first selects q_{ij} to maximize joint surplus, and then negotiates over p_{ij} to split it.

The first-best quantity is determined where the exporter's marginal cost equals the importer's marginal revenue product

$$q_{ij}^* : MR_j(q_{ij}^*) = MC_i(q_{ij}^*).$$

Given q_{ij}^* , the price p_{ij} solves:

$$\max_{p_{ij}} (p_{ij}q_{ij}^* - \Delta C_i(q_{ij}^*))^{1-\phi} (\Delta R_j(q_{ij}^*) - p_{ij}q_{ij}^*)^\phi,$$

where $\Delta R_j(q_{ij}) \equiv R_j(q_j) - \tilde{R}_j(\tilde{q}_j)$ and $\Delta C_i(q_{ij}) \equiv TC_i(q_i) - TC_i(\tilde{q}_i)$ are the incremental revenue for the importer and the incremental cost for the exporter from the match.




FIGURE B.1: Efficient Bargaining

Standard derivations lead to the equilibrium price:

$$p_{ij} = (1 - \phi) \cdot \frac{\Delta R_j(q_{ij})}{q_{ij}} + \phi \cdot \frac{\Delta C_i(q_{ij})}{q_{ij}}$$

a weighted average of the per-unit downstream revenue gain and per-unit upstream cost increase from the match, with bargaining weights given by ϕ . Figure B.1 illustrates this result. The efficient quantity q^* is determined where the exporter's marginal cost $MC(q_i)$ equals the importer's marginal revenue product $p(q_i)$. At q^* , the negotiated price lies between $\Delta C_i/q^* = AC(q^*)$ and $\Delta R_j/q^* = p(q^*)$. The weight ϕ determines the exact price within this interval: as $\phi \rightarrow 0$ the price approaches $\Delta R_j/q^*$, and as $\phi \rightarrow 1$ it approaches $\Delta C_i/q^*$.

B.2 Supply-Driven Quantity Bargaining

In the case of supply-driven quantity bargaining, the exporter first chooses the quantity q_{ij} for a given price p_{ij} to maximize profits, and bargaining occurs over the price, holding the induced supply curve fixed. For tractability, we solve the dual problem where the exporter selects a price p_{ij} for a given quantity q_{ij} , and bargaining takes place over the quantity. Formally, this is expressed as:

$$\begin{cases} \max_{p_{ij}} \pi_i(p_{ij}; q_{ij}) \\ \max_{q_{ij}} [GFT_{ij}^i(q_{ij}, p_{ij})]^{1-\phi} [GFT_{ij}^j(q_{ij}, p_{ij})]^\phi, \end{cases}$$

where the dependence on price and quantity in other parts of the network is left implicit. The solution to the exporter's problem yields the supply function, which we write as:

$$p_{ij}(q_{ij}) = k_i q_i^{\frac{1-\theta}{\theta}}, \quad \text{where } q_i = q_{i(-j)} + q_{ij},$$

as in equation (2.1). Since $p_{ij} = MC_i(q_{ij})$ holds on the supply curve, solving the Nash bargaining problem yields the following first-order condition characterizing the equilibrium quantity (and hence price):

$$\psi_{ij} p_{ij} - MRP_j = \frac{1 - \omega_{ij}^S}{\omega_{ij}^S} (\psi_{ij} - 1) p_{ij},$$

where

$$\psi_{ij} = 1 + c'_{i,q_{ij}} = 1 + \frac{1 - \theta}{\theta} x_{ij},$$




FIGURE B.2: Supply-Driven Quantity Bargaining

and ω_{ij}^S has the same functional form as ω_{ij} in Proposition 1:

$$\omega_{ij}^S \equiv \frac{\phi \psi_{ij} GFT_{ij}^i}{(1 - \phi) GFT_{ij}^j + \phi \psi_{ij} GFT_{ij}^i} \in (0, 1),$$

where $GFT_{ij}^i = p_{ij}q_{ij} - \theta MC_i q_i \Delta_{ij}^x$ and $GFT_{ij}^j = p_j q_j \Delta_{ij}^s - p_{ij}q_{ij}$ are as in Section 2.

The formula has the same structure as ω_{ij} in Proposition 1: ω_{ij}^S is a ratio of the importer's weighted gains $\phi \psi_{ij} GFT_{ij}^i$ to total weighted gains, where the supply-slope factor ψ_{ij} plays the role of $1 + e_i$ (the labor-supply elasticity factor) in their model. The key difference from the baseline is that ω_{ij}^S depends on equilibrium gains-from-trade levels rather than on observable market shares alone, which is why the equilibrium condition does not reduce to a closed-form convex combination.

For $\phi = 0$, $\omega_{ij}^S = 0$: the exporter captures the full importer surplus ($GFT_{ij}^j = 0$), and the equilibrium price is pinned by $\Delta R_j(q_{ij}) = p_{ij} q_{ij}$. For $\phi = 1$, $\omega_{ij}^S = 1$: the first-order condition collapses to $\psi_{ij} p_{ij} = MRP_j$, recovering the standard oligopsony benchmark where the input price is a markdown factor ψ_{ij}^{-1} below the importer's marginal revenue product. Here $\psi_{ij} = 1 +$ the inverse of the importer's residual supply elasticity, reflecting the cost of raising the unit price on all infra-marginal units when the supply curve is upward-sloping.

Figure B.2 illustrates the equilibrium. The blue curve is the exporter's marginal cost MC_i (equivalently, the supply curve), and the black curve is the importer's marginal revenue product MRP_j . The red curve is the importer's marginal factor cost $\psi_{ij} MC_i$, which exceeds MC_i whenever $\theta < 1$. At $\phi = 1$ the equilibrium lies at point A, where $MRP_j = \psi_{ij} MC_i$: the

importer acts as a pure monopsonist, restricting quantity below the efficient level q^* (where $MC_i = MRP_j$). As ϕ decreases, bargaining power shifts toward the exporter, moving the equilibrium up the supply curve toward q^* . The shaded region on the supply curve indicates the feasible price range for $\phi \in (0, 1)$.

A key implication of this protocol is that prices and quantities positively co-move in equilibrium. Because the equilibrium always lies on the exporter's supply curve $p_{ij} = MC_i(q_{ij})$, which is upward-sloping under decreasing returns ($\theta < 1$), any shock that increases the equilibrium quantity also raises the equilibrium price. This stands in sharp contrast to the demand-driven baseline, where the equilibrium lies on the importer's downward-sloping demand schedule and price increases are accompanied by quantity reductions. The data support the demand-driven prediction: tariff-induced price increases are accompanied by quantity declines ($p \uparrow, q \downarrow$).

B.3 Generalized Outside Option

In the baseline model, the importer's (exporter's) gains from trade are the difference between its payoff when trading with all current partners and its payoff when trading with all partners except exporter i (importer j). This rules out two forces that may matter in richer environments: after a failed negotiation, firms may form or reallocate relationships, and broader equilibrium adjustments may shift disagreement payoffs by changing costs or demand elsewhere in the network.

We now consider a more general specification that places less structure on outside options and allows these effects to enter through disagreement payoffs. For tractability, we assume $\varrho = 1$ throughout this section (constant returns to scale for the importer), so that $\eta - 1 = (\nu - 1)\gamma$. Let

$$\Delta_{ij}^{MC_i} = \frac{\widetilde{MC_i}}{MC_i} \quad \text{and} \quad \Delta_{ij}^{c_j} = \frac{\widetilde{MC_j}}{MC_j}$$

denote the percentage change in exporter i 's and importer j 's marginal cost in the event of a failed negotiation. Under this generalization, the gains from trade for firms i and j are given by:

$$\begin{aligned} GFT_{ij}^i(p_{ij}) &= q_{ij} \left(p_{ij} - MC_i \mu_{ij}^{\text{oligopsony}, G} \right), \\ GFT_{ij}^j(p_{ij}) &= \pi_j \Delta_{ij}^{s, G}, \end{aligned}$$

where we define:

$$\mu_{ij}^{\text{oligopsony},G} := \theta \left[\frac{\Delta_{ij}^x}{x_{ij}} \right] \quad \text{with} \quad \Delta_{ij}^x := (1 - \Delta_{ij}^{MC_i}(1 - x_{ij}))$$

and

$$\Delta_{ij}^{s,G} := 1 - (\Delta_{ij}^{c_j})^{1-\nu}.$$

The first-order condition under this generalized outside option implies:

$$\mu_{ij} := \frac{p_{ij}}{MC_i} = (1 - \omega_{ij}^G) \mu_{ij}^{\text{oligopoly}} + \omega_{ij}^G \mu_{ij}^{\text{oligopsony},G}, \quad (\text{B.1})$$

where the weighting factor is given by:

$$\omega_{ij}^G = \frac{\phi \xi^G(GFT_{ij}^j, p_{ij})}{\phi \xi^G(GFT_{ij}^j, p_{ij}) + (1 - \phi)(\varepsilon_{ij} - 1)} \in (0, 1).$$

Here

$$\xi^G(GFT_{ij}^j, p_{ij}) \equiv -\frac{d \ln GFT_{ij}^j}{d \ln p_{ij}} = \frac{(\eta - 1)s_{ij}}{\Delta_{ij}^{s,G}}$$

is the generalized elasticity of the importer's gains from trade with respect to the bilateral price. This expression is equivalent to the baseline elasticity up to the denominator. In the baseline model, that denominator is pinned down by the supplier share alone. Here, instead, it also depends on how the importer's disagreement payoff changes after a failed negotiation, as summarized by $\Delta_{ij}^{s,G}$. Since $\Delta_{ij}^{s,G}$ reflects unobserved changes in marginal cost and outside options, effective bargaining power is no longer identified from s_{ij} alone.

Equation (B.1) shares the weighted-average structure of equation (2.10), but with two key differences. First, the generalized elasticity $\xi^G(GFT_{ij}^j, p_{ij})$ now depends on $\Delta_{ij}^{s,G}$, so effective bargaining power is no longer pinned down by the supplier share alone. Second, the oligopsony markdown depends on Δ_{ij}^x , so it is no longer determined by the buyer share alone.

The main implication is that outside options can absorb both extensive-margin reallocation and broader equilibrium forces, but once they do, the pricing equation depends on unobserved objects. In particular, $\Delta_{ij}^{s,G}$ and Δ_{ij}^x are not recoverable from bilateral market shares alone. Without additional data or stronger assumptions, the model is therefore not identified in a way that is useful for our empirical exercise. The baseline specification imposes more structure precisely to recover markup components from observable bilateral shares and a small set of parameters.

B.4 Full Pass-Through Elasticity

In deriving Proposition 3, we assumed that the shock is applied at the firm-to-firm level and that prices and quantities in other relationships remain fixed. These assumptions allow us to isolate the direct, short-run effects of a shock. However, actual trade policy shocks, such as the Trump tariffs, often apply at the exporter- or product-level and may induce broader adjustments that our static, partial equilibrium model does not capture.

This section extends the analysis to incorporate certain *indirect effects* by capturing how a shock to exporter i influences prices and quantities in other relationships, and how those changes feed back into the bilateral price p_{ij} . While we continue to abstract from full general equilibrium forces, this exercise clarifies how network spillovers may cause reduced-form pass-through estimates to deviate from structural ones, as emphasized by Berger et al. (2022).

We consider an exporter-level price shifter, denoted ϑ_i , and re-derive the relevant elasticities to allow for cross-relationship spillovers. First, the impact on the importer's buyer share becomes:

$$\frac{d \ln x_{ij}}{d \ln \vartheta_i} = -\varepsilon_{ij}(1 - x_{ij}) \frac{d \ln p_{ij}}{d \ln \vartheta_i} + \sum_{z \in \mathcal{Z}_i, z \neq j} x_{iz} \varepsilon_{iz} \frac{d \ln p_{iz}}{d \ln \vartheta_i}.$$

The effect on the exporter's marginal cost is:

$$\frac{d \ln MC_i}{d \ln \vartheta_i} = \frac{1 - \theta}{\theta} \left(-\varepsilon_{ij} x_{ij} \frac{d \ln p_{ij}}{d \ln \vartheta_i} - \sum_{z \in \mathcal{Z}_i, z \neq j} x_{iz} \varepsilon_{iz} \frac{d \ln p_{iz}}{d \ln \vartheta_i} \right).$$

We also account for the fact that p_{ij} affects rival suppliers' shares and prices. The elasticity of supplier share with respect to the shock is:

$$\frac{d \ln s_{ij}}{d \ln \vartheta_i} = (1 - \rho) \frac{d \ln p_{ij}}{d \ln \vartheta_i} \left((1 - s_{ij}) + s_{ij}(1 - \rho) \sum_{k \in \mathcal{Z}_j, k \neq i} s_{kj} \Gamma_{kj}^s \right).$$

Incorporating these indirect effects, the full pass-through elasticity $\Psi_{ij} = \frac{d \ln p_{ij}}{d \ln \vartheta_i}$ is implicitly defined by:

$$\Psi_{ij} = \tilde{\Phi}_{ij} + \tilde{\Phi}_{ij} \left(\Gamma_{ij}^x - \frac{1 - \theta}{\theta} \right) \sum_{z \in \mathcal{Z}_i, z \neq j} x_{iz} \varepsilon_{iz} \Psi_{iz},$$

where:

$$\tilde{\Phi}_{ij} = \left[1 + \Gamma_{ij}^s(\rho - 1) \left((1 - s_{ij}) - s_{ij}(\rho - 1) \sum_{k \in \mathcal{Z}_j, k \neq i} s_{kj} \Gamma_{kj}^s \right) + \Gamma_{ij}^x \varepsilon_{ij} (1 - x_{ij}) + \frac{1 - \theta}{\theta} \varepsilon_{ij} x_{ij} \right]^{-1}.$$

This elasticity Ψ_{ij} embeds two key indirect effects beyond the direct elasticity Φ_{ij} derived earlier.

First, an increase in p_{ij} may cause rivals' supplier shares (e.g., s_{kj}) to rise, increasing their prices p_{kj} via strategic interactions. These adjustments dampen the original substitution away from exporter i , raising $\tilde{\Phi}_{ij}$ relative to Φ_{ij} .

Second, a cost shock to firm i may propagate to other buyers $z \in \mathcal{Z}_i$, affecting p_{iz} , which then feeds back into x_{ij} via firm i 's overall scale and market presence. These changes affect p_{ij} through both the markup and cost channels, amplifying pass-through further.

Together, these network spillovers push the full pass-through elasticity Ψ_{ij} away from the direct elasticity Φ_{ij} . Whether the net effect is amplification or attenuation depends on the strength of substitution patterns and strategic responses, an empirical question. In the main text, we focus on the direct pass-through elasticity Φ_{ij} , which is tightly grounded in our model's Nash-in-Nash structure and match-level pricing assumptions.

B.5 Buyer Power, Decreasing Returns, and Pass-Through

We compare a perfectly competitive benchmark with decreasing returns to scale to our bilateral bargaining framework in order to clarify the source of pass-through heterogeneity. Consider one supplier i and two buyers $j \in \{1, 2\}$. The supplier's marginal cost is

$$MC_i = k_i q_i^{\frac{1-\theta}{\theta}},$$

which implies a marginal cost elasticity

$$mc_q \equiv \frac{d \ln MC_i}{d \ln q_i} = \frac{1 - \theta}{\theta} > 0.$$

Total output is $q_i = q_{i1} + q_{i2}$ and buyer shares are defined as $x_{ij} \equiv q_{ij}/q_i$. For simplicity, both buyers have constant demand elasticity $\sigma \equiv -d \ln q_{ij}/d \ln p_{ij}$.

Under perfect competition and in the absence of price discrimination, prices are uniform and equal to marginal cost, $p_{ij} = p_i = MC_i$. A tariff shock T enters multiplicatively into marginal cost, so that $d \ln p_i = d \ln MC_i + d \ln T$. Using $d \ln MC_i = mc_q d \ln q_i$ together with

$d \ln q_i = -\sigma d \ln p_i$, we obtain

$$d \ln p_i = -mc_q \sigma d \ln p_i + d \ln T,$$

which implies

$$\Phi^{\text{PC}} \equiv \frac{d \ln p_i}{d \ln T} = \frac{1}{1 + mc_q \sigma} = \frac{1}{1 + \frac{1-\theta}{\theta} \sigma}.$$

Pass-through is therefore incomplete due to decreasing returns and identical across buyers.

In the bargaining model, prices are set bilaterally according to $p_{ij} = \mu_{ij} MC_i$. Pass-through becomes relationship-specific:

$$\Phi^{\text{B}} \equiv \frac{d \ln p_{ij}}{d \ln T} = \frac{1}{1 + \Gamma_{ij} + \Lambda_{ij}},$$

where Γ_{ij} captures markup adjustments and

$$\Lambda_{ij} \equiv mc_q x_{ij} \sigma$$

captures the elasticity of marginal cost induced by buyer j 's share of output.

Quantitatively, under our estimate of θ , the markup channel is dominated by the cost channel (see Figure 2), so that $\Gamma_{ij} \approx 0$. Pass-through is then well approximated by

$$\Phi^{\text{B}} \approx \frac{1}{1 + \Lambda_{ij}} = \frac{1}{1 + \frac{1-\theta}{\theta} \sigma x_{ij}}.$$

The two environments share the same underlying mechanism: incomplete pass-through is governed by the slope of marginal cost, mc_q . The difference lies in the unit of adjustment. Under perfect competition, the relevant elasticity is aggregate and pass-through is uniform. Under bargaining, the effective elasticity is share-weighted, generating heterogeneity across relationships. Large buyers internalize a steeper effective supply curve and exhibit lower pass-through, whereas small buyers approximate the constant-returns benchmark. Bargaining thus provides a micro-foundation for decreasing-returns-type pass-through while delivering the relationship-level heterogeneity observed in the data.

C Data Appendix

C.1 Related-Party Trade Measured via Ownership Linkages

A key advantage of the ORBIS dataset is the breadth and detail of its ownership information. It provides comprehensive listings of both direct and indirect shareholders and subsidiaries, along with indicators of each firm’s independence, global ultimate ownership, and group affiliations. This enables us to identify corporate structures at the firm level, including ownership links between firms located in different countries. We define a parent–subsidiary relationship as one in which the parent firm holds at least a 50% ownership stake in the affiliate.

Linking U.S. Importers to Multinational Ownership To identify U.S.-based multinational firms, we match firms in the Census Business Register to their ORBIS counterparts using names, addresses, and GPS coordinates. This linkage combines probabilistic record matching with manual validation, producing a high match rate. As a result, we can flag U.S. establishments that are either majority-owned affiliates of foreign multinationals or parent firms with majority-owned affiliates abroad. This information allows us to identify multinationals with operations in the U.S. without relying solely on the Related Party Trade (RPT) indicator reported in the LFTTD.

Identifying Cross-Border Ownership Links To assess whether the foreign exporter also belongs to the same corporate group, we match the Manufacturer ID (MID) reported in LFTTD to firm records in ORBIS. The MID is constructed by U.S. Customs based on the exporter’s name, address, and country of origin using a set of formatting rules. The MID begins with a two-character country code (or a province code for Canada), followed by a name-based segment derived from the first three letters of the first and second words in the company’s name. If the company name consists of only one word, the first six letters are used. The next segment contains the first four digits from the address number, and the final three characters are the first three alphabetic characters of the city name. Standard formatting conventions apply, including the exclusion of punctuation, one-letter initials, and common stop words such as “the,” “and,” or “of.” Country-specific prefixes (e.g., “OAO” or “ZAO” in Russia or “PT” in Indonesia) are also omitted when constructing the MID.

Using these same rules, we replicate the MID structure for foreign firms in the ORBIS database. We then match each MID in the customs data to candidate firms in ORBIS based on the reconstructed name segment. We assess the quality of each potential match using two dimensions: location and product alignment. A location score is computed based on the

match between city names in the MID and in ORBIS. A product match score is constructed at the MID–firm level by comparing the foreign firm’s NAICS6 industry code in ORBIS to the distribution of HS6 products shipped under that MID in the customs data, using the concordance developed by [Pierce and Schott \(2009\)](#). We retain only those matches where both the location and product match scores exceed 90%. In addition, we drop from the matched LFTTD–ORBIS dataset any ORBIS-linked firm that is associated with fewer than five customs transactions. This filter is based on the number of shipment records matched to each ORBIS identifier in the merged data and is intended to eliminate weak or potentially noisy MID–firm matches.

Another concern is that the MID may sometimes refer to intermediaries rather than manufacturers. Although U.S. customs rules require that the MID correspond to the producer or manufacturer, not to wholesalers or freight forwarders, compliance with this rule is imperfect. To mitigate this concern, we use ORBIS industry codes to exclude retailers, wholesalers, and logistics providers from the matched dataset.

Finally, another challenge with the MID is that it is not a unique firm identifier: a given MID can correspond to multiple legal entities. In our matched data, we address this issue directly by checking whether a MID maps to more than one firm in ORBIS. If multiple firms share the same MID but belong to the same corporate group based on majority ownership links reported in ORBIS, we retain the match. Otherwise, we exclude the ambiguous MID from the analysis.

Taken together, these steps yield a linked dataset that offers a more transparent and conservative definition of related-party trade, based on majority ownership (at least 50%). In contrast to the standard related-party trade (RPT) flag in customs data, which applies a lower threshold of 5% ownership for imports, this approach reduces false positives and more precisely captures transactions where ownership ties are likely to influence pricing. By combining MID-based matching with firm-level ownership structures, the final dataset is well-suited for analyzing pricing behavior in cross-border transactions.

C.2 Data Cleaning, Sample Construction and Summary Statistics

We construct the analysis sample in several steps to align the data with the model and the identification strategy. As in the main text, it is useful to distinguish between restrictions that are essential for identification and those that primarily ensure consistency with the model’s environment.

Preliminary restrictions. We begin by aggregating transaction-level customs records to the annual importer–exporter–product (HS10) level. Because relatively few relationships appear in adjacent months or in the same month across years, annual aggregation provides a natural unit of observation for relationship-level price changes.

To ensure reliable price measurement, we apply several data-quality filters. We drop transactions with missing or zero import value or quantity, invalid exporter identifiers, or U.S. importers that cannot be linked to the Longitudinal Business Database (LBD). We exclude HS codes 98–99 (special provisions), remove observations with log unit values below the 1st or above the 99th percentile of the HS10–country distribution, and trim extreme year-on-year log price changes outside the interval $[-4, 4]$.³⁴ We also exclude energy goods (HS chapter 27), which fall outside the scope of the model.

Identification-driven restrictions. First, we require importer–exporter–product triplets to be observed in at least two consecutive calendar years. This panel structure is necessary to compute year-on-year price changes and to estimate pass-through.

Second, our identification strategy relies on cross-sectional price variation across buyers sourcing the same product from the same supplier. We therefore restrict attention to supplier–product pairs that transact with at least two distinct U.S. buyers in consecutive years. This condition is essential for identifying the bargaining and returns-to-scale parameters, as outlined in Section 4.

Model-alignment restrictions. The remaining restrictions ensure that the empirical setting matches the model’s focus on decentralized bargaining over intermediate inputs. We exclude products classified as consumption goods under the Broad Economic Categories (BEC) system, restricting attention to capital and intermediate inputs.

We also exclude related-party trade, as such transactions may reflect intra-firm transfer pricing rather than arm’s-length bargaining.³⁵ In the baseline, a buyer–supplier pair is classified

³⁴These filters are applied to the estimation sample but retained when computing tenure and relationship length.

³⁵Bernard et al. (2006) document systematic differences between related-party and arm’s-length prices.

as related if ORBIS identifies a common corporate parent. To preserve coverage, we retain all observations not flagged as related by either ORBIS or the LFTTD. For robustness, we consider alternative definitions based on the LFTTD related-party flag and on multinational status (see Appendix C.1).

After applying these restrictions, the final sample isolates repeated, arm’s-length buyer–supplier relationships in intermediate goods with sufficient cross-sectional variation for identification; it retains over 20% of U.S. imports by value and nearly 15% of buyer–supplier–product triplets.

Table C.1 summarizes the cumulative impact of these steps. Panel A focuses on the 2001–2016 sample used for structural estimation. Panel B reports the corresponding summary for the 2017–2018 sample used in the pass-through analysis. The first row of each panel (“All Imports”) includes all U.S. import records for the relevant years. Subsequent rows show the effect of each restriction in turn, including the requirement that buyer–supplier pairs trade the same product in two consecutive calendar years, the exclusion of consumption and energy goods, and the restriction to arm’s-length relationships with sufficient variation for identification.

The final rows report the estimation samples used in the analysis. For 2001–2016 (Panel A), the sample includes approximately \$880 billion in import value, 480,000 buyer–supplier pairs, and 630,000 buyer–supplier–HS10 triplets. For 2017–2018 (Panel B), the corresponding sample comprises \$160 billion in imports, 190,000 pairs, and 250,000 triplets.³⁶ These samples correspond exactly to those used in the structural estimation and post-estimation analysis. Table C.2 reports summary statistics for the main estimation sample, analogous to Table 1 in the main text.

Tables C.3 and C.4 present analogous results for a broader sample that includes capital, intermediate, and consumption goods (“BEC – All (Consec.)”), rather than restricting to non-consumption goods.

³⁶Panel A omits intermediate sample steps and reports only the full and final samples for 2001–2016, as intermediate counts are neither used in the analysis nor releasable under Census disclosure policies. The same applies to Table C.3.

TABLE C.1: Sample Composition by Period – Excluding BEC-Classified Consumption Goods

Sample	Import value (bn USD)	Importers (th)	Exporters (th)	Pairs (th)	Triplets (th)
<i>Panel A: 2001–2016</i>					
All Imports	22,000	1,000	6,900	17,000	40,000
BEC – Non-Cons. (Consec.)	-	-	-	-	-
+ No Energy/Outliers/RPT	-	-	-	-	-
+ Supplier Multi-Buyer	880	70	100	480	630
<i>Panel B: 2017–2018</i>					
All Imports	2,000	330	1,300	2,500	5,300
BEC – All (Consec.)	1,600	160	530	890	1,800
BEC – Non-Cons. (Consec.)	1,000	120	320	540	950
+ No Energy/Outliers/RPT	420	110	270	470	730
+ Supplier Multi-Buyer	160	71	43	190	250

Notes: This table reports sample characteristics for a series of progressively restricted datasets used in the empirical analysis. The first row (“All Imports”) includes all U.S. import records in the sample period. All subsequent rows restrict the sample to buyer–supplier pairs that trade the same HS-10 product in two consecutive calendar years. The “BEC – excl. Cons. (Consec.)” sample includes only capital and intermediate goods, excluding consumption goods as defined by the Broad Economic Categories (BEC) classification. The next sample (“+ No Energy/Outliers”) adds four filters: (i) transactions involving energy-sector goods are excluded; (ii) observations with price levels below the 1st percentile or above the 99th percentile of the within-product price distribution are removed; (iii) extreme log price changes (above 4 or below −4) are excluded; and (iv) related-party transactions, which are defined as trade between entities with ownership ties or corporate control, are dropped following U.S. Census Bureau classification. “+ Supplier Multi-Buyer” restricts to suppliers that trade with at least two different buyers in consecutive years for the same product. “Import value” denotes the total annual value of imports in billions of U.S. dollars. “Importers” and “Exporters” correspond to distinct U.S. buyers and foreign suppliers, respectively. “Pairs” refer to unique buyer–supplier–product matches. “Triplets” refer to unique buyer–supplier–product–year combinations. All figures are reported separately for the 2001–2016 and 2017–2018 periods and are rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. These samples are the exact ones used in the empirical analysis. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE C.2: Summary Statistics for Main Estimation Sample (2001–2018)

Variable	Mean	Std. Dev.	P25	Median	P75
<i>Panel A: Trade Relationships</i>					
s_{ijh} : Supplier share	0.32	0.35	0.03	0.15	0.57
x_{ijh} : Buyer share	0.25	0.29	0.02	0.10	0.40
Relationship length (product h)	4.00	2.80	2.50	3.50	5.50
Relationship length (all products)	4.80	3.30	2.50	4.50	6.50
# Transactions (product h)	120	1100	6.50	16	50
# Transactions (all products)	360	3000	11	36	140
# Products per pair	3.80	7.30	1.50	2.50	4.50
Multi-HS10 dummy	0.59	0.49	0.00	1.00	1.00
# Suppliers per buyer (HS10)	1.80	3.20	1.50	2.50	5.50
Buyer tenure (all products)	9.90	5.00	6.50	10.00	14.00
Buyer tenure (product h)	6.90	4.40	3.50	6.50	10.00
# Buyers per supplier (HS10)	3.20	3.90	2.50	3.50	7.50
Supplier tenure (all products)	8.00	4.60	4.50	8.50	12.00
Supplier tenure (product h)	6.40	4.00	3.50	6.50	9.50
Corr. between s_{ijh} and x_{ijh}	0.041	—	—	—	—
<i>Panel B: Prices (log unit values)</i>					
$\log p$ (pre-duty, excl. charges)	3.40	2.80	1.30	3.00	5.40
$\log p$ (pre-duty)	3.50	2.80	1.40	3.10	5.40
$\log p^{\text{duty}}$ (post-duty)	3.50	2.80	1.40	3.10	5.40

Notes: This table extends Table 1 with additional variables for the main estimation sample at the importer-supplier-product level. Full details on sample construction are provided in Appendix C.2. Panel A reports trade relationship characteristics. s_{ijh} is exporter i 's share in buyer j 's imports of product h ; x_{ijh} is buyer j 's share in exporter i 's U.S. exports of product h . Relationship length (product h) is the number of years between the supplier-buyer pair's first transaction in product h and year t ; Relationship length (all products) is defined analogously, using the pair's first transaction in any product. # Transactions (product h) is the cumulative number of transactions between the pair in product h , from their first transaction through year t ; # Transactions (all products) is defined analogously, summing across all products. # Products per pair is the number of distinct HS10 products traded by the pair in year t , and Multi-HS10 dummy equals one when this number exceeds one. # Suppliers per buyer (HS10) is the number of distinct exporters from which buyer j sources product h in year t ; # Buyers per supplier (HS10) is defined analogously, counting distinct buyers per exporter. Buyer tenure (product h) is the number of years between the buyer's first recorded transaction in product h (with any supplier) and year t ; Buyer tenure (all products) is defined analogously, using the buyer's first recorded transaction in any product. Supplier tenure is defined analogously for the exporter. Corr. between s_{ijh} and x_{ijh} is the sample correlation between the two shares. Panel B reports log unit values (value divided by quantity). $\log p$ (pre-duty, excl. charges) corresponds to the FOB (Free on Board) unit value, excluding shipping and insurance charges; $\log p$ (pre-duty) includes these charges to obtain the CIF (Cost, Insurance, and Freight) unit value; and $\log p^{\text{duty}}$ (post-duty) further incorporates import duties. Statistics are based on confidential LFTTD data and rounded according to U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE C.3: Sample Composition by Period – All BEC Categories

Sample	Import value (bn USD)	Importers (th)	Exporters (th)	Pairs (th)	Triplets (th)
<i>Panel A: 2001–2016</i>					
All Imports	22,000	1,000	6,900	17,000	40,000
BEC – All (Consec.)	-	-	-	-	-
+ No Energy/Outliers/RPT	-	-	-	-	-
+ Supplier Multi-Buyer	1,600	110	210	990	1,500
<i>Panel B: 2017–2018</i>					
All Imports	2,000	330	1,300	2,500	5,300
BEC – All (Consec.)	1,600	160	530	890	1,800
+ No Energy/Outliers/RPT	710	150	470	800	1,500
+ Supplier Multi-Buyer	260	99	79	330	470

Notes: This table reports sample characteristics for a series of progressively restricted datasets used in the empirical analysis. The first row (“All Imports”) includes all U.S. import records in the sample period. All subsequent rows restrict the sample to buyer–supplier pairs that trade the same HS-10 product in two consecutive calendar years. The “BEC – All Categories (Consec.)” sample retains all transactions in capital, intermediate, and consumption goods as defined by the Broad Economic Categories (BEC) system, subject to the consecutive-year condition. The next sample (“+ No Energy/Outliers”) adds four filters: (i) transactions involving energy-sector goods are excluded; (ii) observations with price levels below the 1st percentile or above the 99th percentile of the within-product price distribution are removed; (iii) extreme log price changes (above 4 or below −4) are excluded; and (iv) related-party transactions are dropped following U.S. Census Bureau classification. “+ Supplier Multi-Buyer” restricts to suppliers that trade with at least two different buyers in consecutive years for the same product. “Import value” denotes the total annual value of imports in billions of U.S. dollars. “Importers” and “Exporters” correspond to distinct U.S. buyers and foreign suppliers, respectively. “Pairs” refer to unique buyer–supplier–product matches. “Triplets” refer to unique buyer–supplier–product–year combinations. All figures are reported separately for the 2001–2016 and 2017–2018 periods and are rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE C.4: Summary Statistics – All BEC Categories (2001–2018)

Variable	Mean	Std. Dev.	P25	Median	P75
<i>Panel A: Trade Relationships</i>					
s_{ijh} : Supplier share	0.27	0.32	0.02	0.10	0.43
x_{ijh} : Buyer share	0.27	0.30	0.03	0.13	0.45
Relationship length (product h)	3.90	2.60	2.50	3.50	5.50
Relationship length (all products)	4.60	3.10	2.50	4.50	6.50
# Transactions (product h)	100	890	6.50	16	50
# Transactions (all products)	410	3200	13	45	180
# Products per pair	5.60	12.00	1.50	2.50	5.50
Multi-HS10 dummy	0.68	0.47	0.00	1.00	1.00
# Suppliers per buyer (HS10)	2.00	3.60	1.50	2.50	6.50
Buyer tenure (all products)	9.50	4.90	5.50	9.50	14.00
Buyer tenure (product h)	6.80	4.30	3.50	6.50	10.00
# Buyers per supplier (HS10)	3.00	3.40	2.50	3.50	5.50
Supplier tenure (all products)	7.70	4.40	4.50	7.50	11.00
Supplier tenure (product h)	6.00	3.80	3.50	5.50	9.50
Corr. between s_{ijh} and x_{ijh}	0.053	—	—	—	—
<i>Panel B: Prices (log unit values)</i>					
$\log p$ (pre-duty, excl. charges)	3.30	2.50	1.40	3.00	5.00
$\log p$ (pre-duty)	3.40	2.50	1.50	3.10	5.10
$\log p^{\text{duty}}$ (post-duty)	3.40	2.50	1.50	3.10	5.20

Notes: This table reports summary statistics for the pooled sample, defined analogously to the main estimation sample (Table 1) but without restricting on BEC end-use category, at the importer-supplier-product level. Full details on sample construction are provided in Appendix C.2. Panel A reports trade relationship characteristics. s_{ijh} is exporter i 's share in buyer j 's imports of product h ; x_{ijh} is buyer j 's share in exporter i 's U.S. exports of product h . Relationship length (product h) is the number of years between the supplier-buyer pair's first transaction in product h and year t ; Relationship length (all products) is defined analogously, using the pair's first transaction in any product. # Transactions (product h) is the cumulative number of transactions between the pair in product h , from their first transaction through year t ; # Transactions (all products) is defined analogously, summing across all products. # Products per pair is the number of distinct HS10 products traded by the pair in year t , and Multi-HS10 dummy equals one when this number exceeds one. # Suppliers per buyer (HS10) is the number of distinct exporters from which buyer j sources product h in year t ; # Buyers per supplier (HS10) is defined analogously, counting distinct buyers per exporter. Buyer tenure (product h) is the number of years between the buyer's first recorded transaction in product h (with any supplier) and year t ; Buyer tenure (all products) is defined analogously, using the buyer's first recorded transaction in any product. Supplier tenure is defined analogously for the exporter. Corr. between s_{ijh} and x_{ijh} is the sample correlation between the two shares. Panel B reports log unit values (value divided by quantity). $\log p$ (pre-duty, excl. charges) corresponds to the FOB (Free on Board) unit value, excluding shipping and insurance charges; $\log p$ (pre-duty) includes these charges to obtain the CIF (Cost, Insurance, and Freight) unit value; and $\log p^{\text{duty}}$ (post-duty) further incorporates import duties. Statistics are based on confidential LFTTD data and rounded according to U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

D Additional Empirical Results

D.1 Decomposition of Price Dispersion

To explore the sources of price heterogeneity, we report in Table D.1 the results from OLS regressions decomposing price variation using the specification:

$$\ln p_{ijht} = FE_i + FE_j + FE_{ht} + \beta \mathbf{X}_{ijht} + \varepsilon_{ijht},$$

estimated over the period 2001-2016. We consider three alternative prices: prices that exclude duties ($\ln p_{ijht}$), prices that exclude duties and charges ($\ln p_{ijht}^c$), and prices that include duties and charges ($\ln p_{ijht}^{\text{duty}}$).

Table D.1 finds that controlling for product and year fixed effects explains approximately 50% of the overall price dispersion, while 4% is attributed to match-specific residuals. Notably, this figure changes substantially when isolating variation within supplier-product-year combinations (Panel B), with the buyer-supplier match accounting for 77% of the price variance. This emphasizes that a significant share of price heterogeneity stems from bilateral characteristics that are not solely attributable to either buyer or supplier individually.

TABLE D.1: Fixed-Effect Decomposition of Price Dispersion

Source of Variation	$\ln p_{ijht}$	$\ln p_{ijht}^c$	$\ln p_{ijht}^{\text{duty}}$
<i>Panel A: Overall price dispersion</i>			
FE_{ht}	0.485	0.483	0.486
FE_i	0.424	0.427	0.423
FE_j	0.0464	0.0452	0.0463
Match residual	0.0441	0.0444	0.0442
<i>Panel B: Within exporter–product dispersion</i>			
FE_j	0.233	0.231	0.233
Match residual	0.765	0.768	0.765

Notes: The columns correspond to alternative price definitions: $\ln p_{ijht}$ excludes duties; $\ln p_{ijht}^c$ exclude both duties and charges; $\ln p_{ijht}^{\text{duty}}$ includes duties. The estimation sample includes importer–exporter–product matches observed in two consecutive calendar years, and applies the following restrictions: (i) excludes transactions involving consumption goods (based on the BEC classification), energy-sector products, statistical outliers, and related-party trade; and (ii) retains only suppliers that trade with at least two distinct U.S. buyers in consecutive years. This corresponds to the cumulative sample underlying the “+ Supplier Multi-Buyer” row in Panel B of Table C.1. The control vector $\mathbf{X}_{ijht}$ includes the log of transaction value, the log of relationship longevity (years since the exporter first supplied the buyer with the given HS10 product), and the log of the relative number of partners (the supplier’s number of HS10-level buyers divided by the buyer’s number of HS10-level suppliers). The sample includes 1,200,000 importer–exporter–product–year observations, which have been rounded to four significant digits per U.S. Census Bureau disclosure guidelines. $R^2 = 0.956$. All coefficients in a regression model are significantly different from zero at the 1% significance level. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

D.2 Pass-Through Heterogeneity

Table D.2 examines heterogeneity in tariff pass-through using specification (3.2) with a more demanding set of fixed effects. In addition to product–time fixed effects (FE_{ht}), the specification includes importer–time and exporting country–time fixed effects ($FE_{jt} + FE_{ct}$), accounting for time-varying shocks at the buyer level while flexibly controlling for exporter-side conditions.

Findings replicate the incomplete tariff pass-through estimates reported in Table 3, consistent with exporters adjusting marginal costs in response to demand shifts from large buyers and thereby absorbing a substantial fraction of tariff shocks. In line with Proposition 3, pass-through declines with the buyer share x_{ij} , highlighting the central role of the cost channel.

TABLE D.2: Pass-Through and Relationship Heterogeneity, Stringent Fixed Effects

Dependent variable:	$\Delta \ln p_{ijht}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1 + \tau_{cht})$	-0.184 (0.107)	-0.306 (0.155)	-0.183 (0.093)	-0.113 (0.126)	-0.114 (0.104)	-0.247 (0.145)
$\Delta \ln(1 + \tau_{cht}) \cdot \ln \text{longevity}_{ijht}$		0.088 (0.047)				0.101 (0.047)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$			-0.003 (0.149)		0.004 (0.165)	-0.007 (0.165)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$				-0.282 (0.125)	-0.280 (0.134)	-0.289 (0.128)
$FE_{ht} + FE_{ct} + FE_{jt}$	Yes	Yes	Yes	Yes	Yes	Yes
Observations	249,000	249,000	249,000	249,000	249,000	249,000
R-squared	0.31	0.31	0.31	0.31	0.31	0.31

Notes: This table reports estimates of the pass-through of statutory tariffs, $\Delta \ln(1 + \tau_{cht})$, to duty-exclusive CIF prices at the exporter–importer–product–year level, $\Delta \ln p_{ijht}$. Columns (2) and (6) interact tariffs with the log of relationship longevity, measured as the number of years that buyer j and supplier i have transacted in product h . Columns (3) and (5) interact tariffs with the lagged supplier share, s_{ijht-1} , defined as supplier i 's share in buyer j 's imports of product h . Columns (4) and (5) interact tariffs with the lagged buyer share, x_{ijht-1} , defined as buyer j 's share in supplier i 's exports of product h . All regressions include product–year, exporter country–year, and importer–year fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Controls include: (i) $\ln \text{longevity}_{ijht}$; (ii) $\Delta \ln q_{i(-j)ht}$, the change in exporter i 's total sales of h to U.S. buyers other than j ; and (iii) $\Delta \ln p_{(-i)jht}$, the weighted average price change charged by other suppliers of h to buyer j , using lagged shares as weights. Standard errors are clustered at the HS8 product and exporter-country level. The sample corresponds to the "Supplier Multi-Buyer" definition in Table C.1. Observation counts are rounded to four significant digits per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).


Nonlinear Effects. To explore nonlinearities in tariff pass-through along the distribution of bilateral concentration, we interact the tariff change with quartile dummies of the lagged supplier share ($s_{ijh,t-1}$) and buyer share ($x_{ijh,t-1}$). Specifically, we estimate equation D.1, where $\mathbf{1}\{s_{ijh,t-1} \in Q_q\}$ and $\mathbf{1}\{x_{ijh,t-1} \in Q_q\}$ are indicator variables for quartiles $q = 2, 3, 4$, with the first quartile serving as the omitted category. To separate level and interaction effects, the regression also includes the shares $s_{ijh,t-1}$ and $x_{ijh,t-1}$ themselves as covariates. This specification allows us to test whether pass-through varies nonlinearly across the concentration distribution, while flexibly controlling for underlying differences in market structure. Panel (A) of Figure D.1 includes product–time and exporting country–sector fixed effects, while Panel (B) features product–time, importer–time, and exporting country–time fixed effects.

$$\begin{aligned} \Delta \ln p_{ijht} = & \alpha_0 + \alpha_1 \Delta \ln(1 + \tau_{cht}) + \sum_{q=2}^Q \alpha_{s,q} \cdot \Delta \ln(1 + \tau_{cht}) \cdot \mathbf{1}\{s_{ijh,t-1} \in Q_q\} \\ & + \sum_{q=2}^Q \alpha_{x,q} \cdot \Delta \ln(1 + \tau_{cht}) \cdot \mathbf{1}\{x_{ijh,t-1} \in Q_q\} + \gamma' \mathbf{X}_{ijht} + \mathbf{FE} + \epsilon_{ijht}. \end{aligned} \quad (\text{D.1})$$


The results reveal no evidence of nonlinearities in pass-through with respect to supplier shares (in purple). Interaction coefficients across the upper quartiles of s_{ijht} are uniformly positive, but small and statistically insignificant, indicating that supplier concentration does not materially affect the degree of pass-through. In contrast, buyer shares display a strong monotonic pattern (in green): pass-through declines significantly across higher quartiles of x_{ijht} , consistent with the model’s prediction that dominant buyers limit suppliers’ ability to transmit cost shocks.

FIGURE D.1: Pass-Through by Bilateral Market Share Quartiles

(A) Fixed Effects: $FE_{ht} + FE_{cs}$



(B) Fixed Effects: $FE_{ht} + FE_{ct} + FE_{jt}$



Notes: The figure plots estimated coefficients from regressions of bilateral price changes on tariff changes interacted with quartiles of supplier share (s_{ijht}) and buyer share (x_{ijht}). The first quartile serves as the omitted category. The estimated coefficients correspond to equation (D.1), where $\alpha_{s,q}$ and $\alpha_{x,q}$ capture the interaction of the tariff term with the q th quartile of supplier and buyer shares, respectively. The top panel includes product-year and exporting country-sector fixed effects. The bottom panel includes product-year, importer-year, and exporting country-year fixed effects. Bars indicate 95% confidence intervals. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

Additional Robustness: GE Controls and Price Definitions. Table D.3 presents robustness checks using alternative price definitions and specifications. Columns (1)-(2) exclude the general equilibrium controls included in the baseline model. Columns (3)-(4) use duty-exclusive prices that *exclude* charges, while Columns (5)-(6) use tariff-inclusive prices. In all cases, tariff changes are interacted with lagged supplier and buyer shares, and the set of fixed effects is held constant to ensure comparability across specifications. Estimates remain stable in sign and magnitude across all variants, confirming the robustness of the main results.

TABLE D.3: Additional Robustness: General Equilibrium Controls and Price Definitions

Dependent variable:	$\ln p_{ijht}$ (excl. GE controls)		$\ln p_{ijht}^c$ (before duty, excl. charges)		$\ln p_{ijht}^{\text{duty}}$ (tariff inclusive)	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1 + \tau_{cht})$	-0.025 (0.090)	-0.103 (0.103)	-0.066 (0.093)	-0.163 (0.107)	0.506 (0.099)	0.401 (0.126)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$	0.027 (0.076)	0.011 (0.170)	0.054 (0.072)	0.029 (0.157)	0.017 (0.068)	0.007 (0.143)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$	-0.403 (0.112)	-0.287 (0.135)	-0.403 (0.113)	-0.271 (0.135)	-0.421 (0.121)	-0.287 (0.145)
$FE_{ht} + FE_{cs}$	Yes	No	Yes	No	Yes	No
$FE_{ht} + FE_{ct} + FE_{jt}$	No	Yes	No	Yes	No	Yes
Observations	249,000	249,000	249,000	249,000	249,000	249,000
R-squared	0.04	0.31	0.04	0.31	0.05	0.31

Notes: This table reports robustness checks on tariff pass-through specifications using alternative price definitions and control sets. Columns (1)–(2) exclude general equilibrium controls; Columns (3)–(4) use pre-duty prices excluding charges; and Columns (5)–(6) use tariff-inclusive prices. In each case, we report specifications using either baseline fixed effects ($FE_{ht} + FE_{cs}$) or a more stringent set of fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Standard errors are clustered at the HS8 product and exporter-country level. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

D.3 Additional Results on Model Estimation

We assess the robustness of our structural estimates in Section 4.2. We first re-estimate the main GMM specification using an expanded sample that includes all Broad Economic Categories (BECs). We also estimate the structural parameters on a sample restricted to consumption goods only. Results are reported in Tables D.4 and D.5, respectively

Estimated returns to scale are systematically higher for consumption goods. Restricting the estimation to consumption goods raises $\hat{\theta}$ from 0.42–0.51 in the baseline to 0.54–0.66. Expanding the sample to include all BEC categories also increases estimates, with $\hat{\theta}$ ranging from 0.48 to 0.60, consistent with the higher values observed for consumption goods. This pattern suggests that consumption goods are less relationship-specific and therefore face more elastic supply conditions, implying greater substitutability across suppliers and higher estimated returns to scale in the model.

Finally, Table D.6 examines the robustness of our structural estimates to an alternative definition of related-party trade. Instead of using our baseline ORBIS-based classification, we re-estimate the model excluding transactions identified as related-party trade using the standard RPT indicator in the U.S. Census data. Appendix C.1 explains the difference between the two indicators. Reassuringly, the estimated bargaining-power and returns-to-scale parameters remain very close to the baseline results, indicating that our main conclusions are not driven by the particular way related-party relationships are classified.

TABLE D.4: Estimated Model Primitives - All BEC Categories

Panel A: Calibrated Parameters				
$\hat{\nu}$	$\hat{\gamma}$		$\hat{\rho}$	
4	0.5		10	
Panel B: Estimated Parameters (GMM)				
	(1)	(2)	(3)	(4)
Rel. bargaining power: $\ln \widehat{\frac{\phi}{1-\phi}}$	1.180 (0.023)		0.940 (0.021)	
Returns to scale ($\hat{\theta}$)	0.508 (0.002)	0.585 (0.003)	0.483 (0.003)	0.595 (0.003)
Constant		3.355 (0.196)		1.299 (0.103)
Longevity		0.115 (0.026)		0.902 (0.090)
Number of HS10 transactions		-0.315 (0.021)		-0.090 (0.016)
Multiple HS10 dummy		0.137 (0.028)		0.199 (0.029)
Lagged outside option		-0.267 (0.019)		-0.272 (0.023)
None	Yes	Yes	No	No
$FE_h + FE_t + FE_j$	No	No	Yes	Yes
Observations	6,143,000			
Panel C: Implied Bargaining Powers ($\hat{\phi}$)				
Mean	0.765 (0.004)	0.923 (0.075)	0.719 (0.004)	0.899 (0.095)
Median	— —	0.946 (0.075)	— —	0.928 (0.095)

Notes: This table presents model estimates based on a sample that includes all Broad Economic Categories (BEC), including consumption goods, for the period 2001-2016. Panel A reports calibrated parameters: demand elasticity (ν), cost elasticity to foreign input prices (γ), and the elasticity of substitution across foreign varieties (ρ). Panel B reports GMM estimates. Columns (1) and (3) impose a constant ϕ , while Columns (2) and (4) allow heterogeneous bargaining power via κ . Specifications differ in fixed effects. Controls include relationship longevity, transaction intensity, the relative outside option (lagged), and a multi-product indicator. Panel C reports the mean and median implied bargaining power. Robust standard errors; Panel C uses the delta method. Instruments include HS10-level counts of exporters and importers and lagged bilateral shares (excluding the focal pair). All estimations use CIF duty-exclusive unit values. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE D.5: Estimated Model Primitives - BEC Consumption Categories

Panel A: Calibrated Parameters				
$\hat{\nu}$	$\hat{\gamma}$		$\hat{\rho}$	
4	0.5		10	
Panel B: Estimated Parameters (GMM)				
	(1)	(2)	(3)	(4)
Rel. bargaining power: $\ln \widehat{\frac{\phi}{1-\phi}}$	0.936 (0.023)		0.729 (0.021)	
Returns to scale ($\hat{\theta}$)	0.558 (0.003)	0.646 (0.004)	0.539 (0.003)	0.664 (0.004)
Constant		2.501 (0.213)		0.568 (0.093)
Longevity		0.628 (0.070)		1.414 (0.198)
Number of HS10 transactions		-0.386 (0.037)		-0.137 (0.030)
Multiple HS10 dummy		0.452 (0.056)		0.344 (0.052)
Lagged outside option		-0.266 (0.027)		-0.299 (0.038)
None	Yes	Yes	No	No
$FE_h + FE_t + FE_j$	No	No	Yes	Yes
Observations	3,024,000			
Panel C: Implied Bargaining Powers ($\hat{\phi}$)				
Mean	0.718 (0.005)	0.919 (0.077)	0.674 (0.005)	0.891 (0.111)
Median	— —	0.942 (0.077)	— —	0.928 (0.111)

Notes: This table presents model estimates based on a sample that only includes consumption goods based on Broad Economic Categories (BEC), for the period 2001-2016. Panel A reports calibrated parameters: demand elasticity (ν), cost elasticity to foreign input prices (γ), and the elasticity of substitution across foreign varieties (ρ). Panel B reports GMM estimates. Columns (1) and (3) impose a constant ϕ , while Columns (2) and (4) allow heterogeneous bargaining power via κ . Specifications differ in fixed effects. Controls include relationship longevity, transaction intensity, the relative outside option (lagged), and a multi-product indicator. Panel C reports the mean and median implied bargaining power. Robust standard errors; Panel C uses the delta method. Instruments include HS10-level counts of exporters and importers and lagged bilateral shares (excluding the focal pair). All estimations use CIF duty-exclusive unit values. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE D.6: Estimated Model Primitives under Alternative Related-Party Classification

	(1)	(2)	(3)	(4)
Implied Mean Bargaining Power ($\hat{\phi}$)	0.837 (0.008)	0.928 (0.078)	0.761 (0.008)	0.899 (0.082)
Returns to scale ($\hat{\theta}$)	0.455 (0.004)	0.500 (0.005)	0.410 (0.005)	0.504 (0.006)
None	Yes	Yes	No	No
$FE_h + FE_t + FE_j$	No	No	Yes	Yes

Notes: Model estimates are based on U.S. arm’s-length imports of intermediate inputs and capital goods over 2001–2016, excluding transactions identified as related-party trade using the RPT indicator in the U.S. Census data; see Appendix C.1 for details. Columns (1) and (3) impose a constant ϕ , while Columns (2) and (4) allow for heterogeneous bargaining power through κ , parameterized as a function of relationship longevity, transaction intensity, the lagged relative outside option, and a multi-product indicator. The demand elasticity is set to $\nu = 4$, the cost elasticity to foreign input prices to $\gamma = 0.5$, and the elasticity of substitution across foreign varieties to $\rho = 10$. Specifications differ in the fixed effects included. Robust standard errors are reported in parentheses. All estimations use CIF duty-exclusive unit values. All regressions include 3,120,000 observations. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

D.4 Additional Results on Model Fit

Pass-through and Relationship Heterogeneity: Data vs. Model. Table D.7 reports tariff pass-through in the data and in the model (baseline and restricted variants) under a more demanding set of fixed effects. Panel A shows that the model captures the main pass-through effect in the data very closely. Panel B shows how pass-through varies with buyer and supplier concentration, measured by lagged buyer and supplier shares. The baseline model closely replicates these patterns: pass-through declines sharply with buyer concentration (x_{ijht-1}) and increases modestly with supplier concentration (s_{ijht-1}), with magnitudes comparable to those observed in the data.

IV-Based Goodness-of-Fit Test. We formally assess model fit using an IV regression following Adão et al. (2023):

$$\Delta \ln p_{ijht} = \beta \widehat{\Delta \ln p_{ijht}} + \mathbf{FE} + u_{ijht}, \quad (\text{D.2})$$

where $\Delta \ln p_{ijht}$ is the observed change in the duty-inclusive price and $\widehat{\Delta \ln p_{ijht}}$ is the corresponding model-implied change. Because observed price changes may reflect shocks unrelated to tariffs, we instrument $\widehat{\Delta \ln p_{ijht}}$ with statutory tariff changes to isolate tariff-driven variation.³⁷ Under the null that the model correctly captures tariff-induced price movements, the IV coefficient $\hat{\beta}$ equals one.

Figure D.2 reports the IV estimates of β across the four model variants under the two fixed-effects specifications. Blue markers correspond to regressions with product–time and country–sector fixed effects, while red markers use product–time, country–time, and buyer–time fixed effects. The model generates price changes closely aligned with those observed in the data. Formally, however, we cannot reject any of the variants at conventional levels of significance.

Sensitivity of Model Fit to Alternative Parameter Values. In our baseline estimation we set $\varrho = 1$, $\nu = 4$, and $\gamma = 0.5$. As all these parameters enter the model only through η , fixing the buyer’s returns to scale parameter ϱ to one does not restrict generality. Nevertheless, we conduct a robustness exercise where we set $\varrho = 0.5$, close to the returns to scale parameter we estimate for the supplier side. Further, we also consider a value of $\rho = 5$, the lower end of the range provided in Anderson and van Wincoop (2004).

³⁷Unlike the pass-through analysis in the main text, which uses duty-exclusive prices to mitigate measurement error in applied duties, the IV strategy isolates variation induced by statutory tariffs and evaluates the model against duty-inclusive price changes.

Table D.8 presents IV-based goodness-of-fit tests as in equation (D.2) for these alternative set of calibrated parameters of the model and two alternative sets of fixed effects. The results show that models incorporating bargaining and decreasing returns to scale (Columns (1) and (5)) have coefficients closer to one, indicating empirical alignment, consistent with the results in Figure D.2.

Model Fit: Quantity and Sales Responses. Table D.9 compares how quantities respond to tariffs in the data (Panel A) and in the model (Panel B). Columns (1) and (2) use baseline fixed effects; Columns (3) and (4) add more demanding fixed effects. Columns (2) and (4) include interactions with supplier and buyer shares.

The table shows that the model generates sizable average quantity declines and predicts heterogeneity across relationships. In contrast, the interaction terms with supplier and buyer shares for the quantity responses in the data are statistically imprecise, suggesting inconclusive evidence of heterogeneity in the data.

Table D.10 reports IV-based tests comparing observed quantity (Panel A) and sales (Panel B) changes to those implied by the model under alternative parameterizations. The model predicts

$$\widehat{\Delta \ln q_{ijht}} = -\varepsilon_{ijht} \widehat{\Delta \ln p_{ijht}},$$

where ε_{ijht} denotes the match-specific demand elasticity and $\widehat{\Delta \ln p_{ijht}}$ is defined in equation (D.2). For sales changes, we use:

$$\widehat{\Delta \ln r_{ijht}} = (1 - \varepsilon_{ijht}) \widehat{\Delta \ln p_{ijht}}.$$


Although the formal tests reject all specifications, the specification that combines bargaining with decreasing returns and the specification with decreasing returns only provide the best quantitative fit, consistent with the results for prices.

TABLE D.7: Price Responses and Relationship Heterogeneity: Data vs. Model

Panel A: Main Effect (Stringent Fixed Effects)					
	Data	Baseline	$\phi = 0,$ $\theta = 1$	$\theta = 1$	$\phi = 0$
$\Delta \ln(1 + \tau_{cht})$	-0.184 (0.107)	-0.226 (0.009)	-0.095 (0.008)	-0.014 (0.001)	-0.286 (0.011)
R-squared	0.31	0.48	0.42	0.41	0.51
Panel B: Including Interactions (Stringent Fixed Effects)					
	Data	Baseline	$\phi = 0,$ $\theta = 1$	$\theta = 1$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln(1 + \tau_{cht})$	-0.114 (0.104)	-0.137 (0.011)	-0.030 (0.007)	-0.002 (0.001)	-0.200 (0.014)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$	0.004 (0.165)	0.116 (0.013)	-0.120 (0.022)	-0.022 (0.003)	0.073 (0.013)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$	-0.280 (0.134)	-0.433 (0.027)	-0.078 (0.020)	-0.013 (0.004)	-0.372 (0.025)
R-squared	0.31	0.57	0.45	0.45	0.56
$FE_{ht} + FE_{ct} + FE_{jt}$	Yes	Yes	Yes	Yes	Yes
Observations	249,000				

Notes: This table reports the response of duty-exclusive CIF prices to tariff changes at the exporter–importer–product level. Column (1) presents reduced-form estimates using observed price changes in the data. Columns (2)–(5) report the corresponding responses implied by model-predicted price changes under each specification. Panel A reports the main effect of tariff changes, while Panel B allows the response to vary with lagged supplier shares (s_{ijht-1}) and buyer shares (x_{ijht-1}). All specifications include product–time, country–time, and buyer–time fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Standard errors are clustered at the HS8 product and exporter–country level. Observation counts are rounded to four significant digits per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

FIGURE D.2: IV-Based Goodness-of-Fit Test



Notes: Each point reports the coefficient from an IV regression of observed log duty-inclusive price changes on model-predicted changes $\Delta \ln \widehat{p}_{ijht}$, using statutory tariffs as instruments. Lines show 95% confidence intervals. Blue and red denote regressions with $FE_{ht} + FE_{cs}$ and $FE_{ht} + FE_{ct} + FE_{jt}$, respectively. Standard errors are clustered by product and exporter-country. Observation counts (249,000) are rounded per Census disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE D.8: IV-Based Goodness-of-Fit Test, Alternative Parameters

Panel A: Baseline Fixed Effects

Dependent variable:	$\Delta \ln p_{ijht}$							
	$\rho = 5$				$\rho = 0.5$			
	Baseline $\phi = 0,$ $\theta = 1$	$\phi = 0,$ $\theta = 1$	$\phi = 0$	Baseline	$\phi = 0,$ $\theta = 1$	$\phi = 0,$ $\theta = 1$	$\phi = 0$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\Delta \ln p_{ijht}}$	1.158 (0.331)	0.889 (0.253)	0.771 (0.220)	1.391 (0.397)	1.291 (0.368)	0.936 (0.267)	0.776 (0.221)	1.619 (0.4763)
$FE_{ht} + FE_{cs}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Stringent Fixed Effects

Dependent variable:	$\Delta \ln p_{ijht}$							
	$\rho = 5$				$\rho = 0.5$			
	Baseline $\phi = 0,$ $\theta = 1$	$\phi = 0,$ $\theta = 1$	$\phi = 0$	Baseline	$\phi = 0,$ $\theta = 1$	$\phi = 0,$ $\theta = 1$	$\phi = 0$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\Delta \ln p_{ijht}}$	0.998 (0.341)	0.776 (0.265)	0.667 (0.227)	1.215 (0.415)	1.128 (0.386)	0.805 (0.275)	0.670 (0.229)	1.432 (0.490)
$FE_{ht} + FE_{ct} + FE_{jt}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	249,000							

Notes: Each column reports the coefficient from an IV regression of the observed change in log duty-inclusive price on the corresponding model-predicted change $\widehat{\Delta \ln p_{ijht}}$, using statutory tariff changes as instruments. Columns (1)-(4) use predicted prices using a $\rho = 5$ instead of $\rho = 10$ as in the baseline. Columns (5)-(8) use predicted prices using a $\rho = 0.5$ instead of $\rho = 1$ as in the baseline. Notice that in each case, we re-estimate the calibrated parameters ϕ and θ , reported in Appendix D.3 In Panel A, all columns include product-time and country-sector fixed effects ($FE_{ht} + FE_{cs}$). In Panel B, all columns include product-time, country-time, and buyer-time fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Standard errors are clustered at the product and exporter-country level. The number of observations is rounded to four significant digits in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE D.9: *Quantity* Responses and Relationship Heterogeneity: Data vs. Model

	Baseline FE		Stringent FE	
	Data (1)	Model (2)	Data (3)	Model (4)
<i>Panel A: Main effects</i>				
$\Delta \ln(1 + \tau_{cht})$	-0.568 (0.249)	-1.710 (0.075)	-1.021 (0.185)	-1.921 (0.106)
R-squared	0.06	0.25	0.32	0.40
<i>Panel B: Including Interactions</i>				
$\Delta \ln(1 + \tau_{cht})$	0.006 (0.331)	-3.683 (0.203)	-0.761 (0.237)	-3.705 (0.229)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$	-0.511 (0.323)	2.755 (0.461)	-0.278 (0.276)	2.816 (0.617)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$	-0.201 (0.229)	2.453 (0.242)	-0.130 (0.242)	2.703 (0.349)
R-squared	0.10	0.36	0.36	0.49
$FE_{ht} + FE_{cs}$	Yes	Yes	No	No
$FE_{ht} + FE_{ct} + FE_{jt}$	No	No	Yes	Yes
Observations	249,000			

Notes: Panel A reports main effects of tariff changes on quantities. Panel B introduces interactions between tariff changes and lagged supplier share (s_{ijht-1}) and lagged buyer share (x_{ijht-1}). Columns (1) and (2) use baseline fixed effects ($FE_{ht} + FE_{cs}$), while Columns (3) and (4) use a more demanding specification with product-year, country-year, and buyer-year fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Standard errors are clustered at the HS8 product and exporter-country level. Observation counts are rounded to four significant digits per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE D.10: IV-Based Goodness-of-Fit Test for Quantities and Sales

<i>Panel A: Quantities</i>								
Dependent Variable:	$\Delta \ln q_{ijht}$							
	Baseline		$\phi = 0, \theta = 1$		$\theta = 1$		$\phi = 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\Delta \ln q_{ijht}}$	0.401 (0.165)	0.592 (0.161)	0.242 (0.100)	0.361 (0.097)	0.213 (0.087)	0.320 (0.086)	0.496 (0.205)	0.726 (0.198)
<i>Panel B: Sales</i>								
Dependent Variable:	$\Delta \ln(p_{ijht} \cdot q_{ijht})$							
	Baseline		$\phi = 0, \theta = 1$		$\theta = 1$		$\phi = 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\widehat{\Delta \ln(p \cdot q)_{ijht}}$	0.206 (0.184)	0.489 (0.177)	0.121 (0.108)	0.292 (0.105)	0.107 (0.095)	0.262 (0.094)	0.255 (0.228)	0.598 (0.216)
$FE_{ht} + FE_{cs}$	Yes	No	Yes	No	Yes	No	Yes	No
$FE_{ht} + FE_{ct} + FE_{jt}$	No	Yes	No	Yes	No	Yes	No	Yes
Observations	249,000							

Notes: Each column reports the coefficient from an IV regression of the observed change in log quantity (Panel A) or log sales (Panel B) on the corresponding model-predicted change, using statutory tariff changes as instruments. Columns (1), (3), (5), and (7) include product–time and country–sector fixed effects ($FE_{ht} + FE_{cs}$), while Columns (2), (4), (6), and (8) include product–time, country–time, and buyer–time fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Standard errors are clustered at the HS8 product and exporter–country level. Observation counts are rounded to four significant digits per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

D.5 Additional Pass-Through Results

D.5.1 Pass-Through Across Samples

Table D.11 examines how tariff pass-through estimates vary across alternative sample restrictions under baseline fixed effects (product–time and country–sector; Panel A) and more demanding specifications that additionally include country-time (instead of country-sector) and buyer–time fixed effects (Panel B). For reference, the first two columns replicate Columns (1) and (2) of Table 6. These estimates imply pass-through of roughly 82–88 percent in the data, closely aligned with the model’s predictions, with no statistically significant differences.

As the sample is progressively broadened—from including final goods (column 3), to incorporating single-buyer matches (column 4), and finally to the most inclusive specification, which adds related-party transactions, energy goods, and outliers (column 5)—estimated pass-through rises monotonically, reaching 94–96 percent.

These results indicate that sample composition plays an important role in measured pass-through. At the same time, when attention is restricted to stable buyer–supplier relationships, pass-through remains incomplete and below the near-100 percent estimates often reported in the literature. To further reconcile our findings with this evidence, we next aggregate the data to the product level and exploit monthly variation.

D.5.2 Product-Level Pass-Through Estimates

To complement our main analysis, we replicate a standard pass-through specification using data aggregated at the product-country-month level. We construct the data directly from the buyer–supplier–product triplets that form the basis of our firm-level regressions.

While much of the literature analyzes monthly price and tariff changes at the product–country level, such approaches reflect both intensive and extensive margin adjustments, including changes in trading partners or the entry and exit of relationships. In contrast, our aggregation focuses exclusively on consecutive transactions between the same buyer and supplier for a given product. This setup isolates price responses within ongoing relationships, capturing what is arguably the most direct expression of tariff pass-through at the micro level.

Data We measure price changes at the HS10–source-country–month level using the same set of buyer–supplier–product links as in the baseline analysis, restricted to relationships with at least one transaction in both 2017 and 2018. For each relationship, we retain the month in which transactions occur in both years. We then aggregate values and quantities across these transactions to the product–source-country–month level, construct unit values,

and compute month-to-month price changes.

Relative to Table D.11, this procedure introduces two changes: prices are aggregated across trading relationships to the product–source–country level, and price changes are measured at the monthly frequency rather than annually. The objective of Table D.12 is to evaluate how pass-through estimates change when prices are aggregated across firms within a product–source–country–month cell, while maintaining the same underlying set of buyer–supplier links observed in both 2017 and 2018.

Column (1) reports results for the baseline sample. The sample is then progressively broadened along the same margins as in Table D.11: Column (2) adds consumption goods, Column (3) allows suppliers serving only one buyer, Column (4) further includes related-party transactions, energy goods, and extreme price changes, and Column (5) additionally incorporates non-consecutive buyer–supplier relationships.

Results Across specifications, we continue to find incomplete tariff pass-through to U.S. import prices. When the aggregation is performed using the same set of repeated buyer–supplier relationships as in the baseline analysis, pass-through remains well below one, consistent with the match-level evidence. As the sample is progressively broadened to include consumption goods, single-buyer suppliers, related-party transactions, energy goods, and eventually non-consecutive relationships, estimated pass-through increases steadily, approaching the near-complete transmission typically found in the literature.

Differences between match-level and product-level pass-through therefore primarily reflect the set of relationships included in the data rather than aggregation itself. Incomplete pass-through remains a robust feature when attention is restricted to persistent buyer–supplier links.

TABLE D.11: Tariff Pass-Through Across Different Samples

Dependent variable: $\Delta \ln p_{ijht}$	Model	Baseline	+ Final Goods	+Suppliers w/ < 2 Buyers	+ Energy/ RPT/ Outliers
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Baseline Fixed Effects</i>					
$\Delta \ln(1 + \tau_{cht})$	-0.222 (0.008)	-0.126 (0.091)	-0.151 (0.089)	-0.138 (0.034)	-0.060 (0.044)
$FE_{ht} + FE_{cs}$	Yes	Yes	Yes	Yes	Yes
R-squared	0.31	0.04	0.04	0.02	0.02
<i>Panel B: Stringent Fixed Effects</i>					
$\Delta \ln(1 + \tau_{cht})$	-0.226 (0.009)	-0.184 (0.107)	-0.118 (0.075)	-0.083 (0.032)	-0.036 (0.042)
$FE_{ht} + FE_{ct} + FE_{jt}$	Yes	Yes	Yes	Yes	Yes
R-squared	0.48	0.31	0.26	0.14	0.12
Observations	249,000	249,000	473,000	1,490,000	1,768,000

Notes: This table reports tariff pass-through estimates to duty-exclusive CIF prices at the exporter-importer-product level. Panel A uses baseline fixed effects: product-time and country-sector ($FE_{ht} + FE_{cs}$). Panel B uses a more stringent specification: product-time, country-time, and buyer-time ($FE_{ht} + FE_{ct} + FE_{jt}$). Column (1) uses the model-predicted price change as the dependent variable. Column (2) uses the observed price change in the baseline sample and is identical to Column (1) of Table 3. Column (3) adds relationships in which the supplier trades with only one U.S. importer. Column (4) further expands the sample to include relationships that are either related parties, involve energy commodities, or exhibit extreme price levels or changes. Column (5) incorporates consumption goods, thus encompassing all consecutive exporter-importer-product combinations. Standard errors are clustered at the HS8 product and country level. Observation counts are rounded per U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

TABLE D.12: Product-Level Tariff Pass-Through Across Alternative Samples

Dependent variable:	Baseline	+ Final	+ Suppliers	+ Energy/	+ Non-Consecutive
$\Delta \ln p_{ijht}$		Goods	w/ < 2 Buyers	RPT/ Outliers	Relationships
	(1)	(2)	(3)	(4)	(5)
$\Delta \ln(1 + \tau_{cht})$	-0.654 (0.245)	-0.353 (0.156)	-0.063 (0.073)	-0.054 (0.068)	-0.017 (0.068)
$FE_{ht} + FE_{cs}$	Yes	Yes	Yes	Yes	Yes
R-squared	0.23	0.22	0.14	0.13	0.11
Observations	124,000	213,000	665,000	805,000	1,071,000

Notes: This table reports product-level pass-through estimates. Product-level observations are constructed from firm-buyer-HS relationships that transact in both years t and $t-1$. Values and quantities are aggregated to the product-source-country-month level, unit values are constructed at that frequency, and month-to-month changes in log CIF prices are computed. Column (1) reports the baseline sample, restricted to repeated arm's-length relationships in intermediate goods among multi-buyer suppliers. Column (2) broadens the sample to include consumption goods (all BEC categories). Column (3) further adds suppliers that trade with only one U.S. importer. Column (4) further includes related-party transactions, energy commodities, and observations with extreme price levels or changes. Column (5) removes the consecutive-year restriction and includes all exporter-importer-product combinations. Standard errors are clustered at the HS8 product and country level. Observation counts are rounded in accordance with U.S. Census Bureau disclosure guidelines. *Source:* FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).


E Estimation Appendix

E.1 Calibration of Downstream Demand Elasticity (ν)

Consider an extension to our model in which importer j sells its output q_j to downstream consumers in different countries. In each country, a representative consumer allocates expenditure between domestic and imported goods, and the sub-utility over imported varieties is CES with elasticity of substitution σ_g at the HS10 good level. Broda and Weinstein (2006) (BW) estimate these import demand elasticities for the United States.

In our framework, the parameter ν governs the own-price elasticity of demand faced by importer j 's output in the downstream market. We interpret BW's σ_g estimates as measuring this variety-level demand elasticity. Accordingly, we calibrate ν using their estimates. The distribution of σ_g is shown in Figure E.1. Given a mean value of 3.85 and a median of 2.8, we set $\nu = 4$, which we view as a conservative choice.

FIGURE E.1: Downstream Demand Elasticity



Notes: The figure displays the estimates of the import demand elasticity σ_g from Broda and Weinstein (2006). The mean and median value of σ_g^{US} are 3.85 and 2.8, respectively. Estimates are truncated above at 20, and below at 1.

E.2 Potential Endogeneity from Linkage Formation

The marginal cost, from equation (2.1), can be written as

$$\ln MC_{iht} = \frac{1-\theta}{\theta} \ln q_{iht} + \ln k_{iht}.$$

As noted in the main text, we generalize k_{iht} to allow for match-specific cost components and denote as k_{ijht} . This buyer–supplier–time specific term captures heterogeneity in effective marginal cost arising from match-specific productivity, quality differences, or other unobserved relationship-level factors.

Accordingly, we interpret $\ln(k_{ijht})$ as a structural component of marginal cost, rather than an econometric residual. Although we do not explicitly model the determinants of k_{ijht} , in estimation we treat it nonparametrically as a buyer–supplier–product–time varying term. This approach captures idiosyncratic price differences not explained by bilateral markups and mitigates concerns related to selection into matches.

Here we discuss how selection operating through k_{ijht} could affect identification of the structural parameters. We model the value of a potential match between exporter i and importer j as

$$V(X_{ijht}) = f(X_{ijht}) + \Xi_{ijht},$$

where X_{ijht} denotes observed match characteristics, $f(\cdot)$ is a deterministic function, and Ξ_{ijht} captures unobserved, idiosyncratic variation. A match is observed if and only if $V(X_{ijht}) > 0$, which implies $\Xi_{ijht} > -f(X_{ijht})$.

Hence, the observed transaction (q_{ijht}, p_{ijht}) satisfies

$$q_{ijht}^* = \begin{cases} 0, & \text{if } \Xi_{ijht} \leq -f(X_{ijht}), \\ q_{ijht}(p_{ijht}), & \text{otherwise,} \end{cases} \quad (\text{E.1})$$

where $q_{ijht}(\cdot)$ corresponds to the demand function defined in equation (2.6). When $\Xi_{ijht} \leq -f(X_{ijht})$, no match is formed, and bilateral prices and quantities are not observed.

This framework clarifies how selection may bias estimation. If the unobserved component Ξ_{ijht} is correlated with key determinants of bilateral markups (s_{ijht} and x_{ijht}) then selection induces endogeneity by generating correlation between the regression error and the regressors. In that case, estimates of ϕ and θ may be biased, with the sign and magnitude of the bias depending on the underlying correlation structure.

Even if Ξ_{ijht} is uncorrelated with the regressors in the population, selection based on Ξ_{ijht}

can induce correlation between the regression error and the regressors whenever Ξ_{ijht} is correlated with the selection rule k_{ijht} , a standard truncation problem. The absence of bias therefore requires the stronger condition, that is, selection must be random conditional on observables.

In practice, if the unobserved match surplus Ξ_{ijht} primarily reflects exporter-specific fundamentals or shocks that are stable across a supplier’s buyers, differencing across buyers within the same exporter mitigates this source of endogeneity. Our estimation strategy exploits cross-sectional variation in prices across buyers of the same supplier and is designed to reduce, though not eliminate, concerns about selection bias.

When $\text{Cor}(k_{ijht}, \Xi_{ijht}) \neq 0$, and k_{ijht} is also correlated with the bilateral market shares that determine the markup, the resulting bias can be positive or negative depending on the sign and strength of these correlations. However, if the unobserved match surplus Ξ_{ijht} reflects exporter-specific fundamentals or shocks that are relatively stable across a supplier’s buyers, then differencing across matches within the same exporter helps mitigate this source of endogeneity. Our estimation strategy, which leverages cross-sectional variation in prices across buyers of the same supplier, is therefore designed to partially address this endogeneity concern.

We further address the potential endogeneity of k_{ijht} using instrumental variables. Specifically, we instrument for bilateral shares using the number of importers and exporters active in product h , which proxy for the broader population of potential foreign suppliers and U.S. buyers. These variables are plausibly exogenous to the match-specific surplus Ξ_{ijht} , and are not determined by the selection equation (E.1) above. The exclusion restriction is that Ξ_{ijht} is uncorrelated with the overall size of the trading population in product h . Intuitively, a larger number of potential partners shifts the distribution of the order statistics of Ξ_{ijht} , influencing the likelihood of match formation without directly affecting whether a given match exceeds the threshold $-f(X_{ijht})$.

We acknowledge that alternative network formation processes could give rise to different endogeneity concerns. For instance, in dynamic formation games, k_{ijht} may correlate with the expected continuation value of future matches, which would introduce additional selection channels. In such cases, the direction of the bias on ϕ and θ is harder to characterize. Nonetheless, the static selection framework outlined above provides a useful benchmark that aligns with our theoretical setup and clarifies the potential sources of endogeneity relevant for our identification strategy.