Two-Sided Market Power in Firm-to-Firm Trade*

Alviarez, V. — IADB
Fioretti, M. — Bocconi
Kikkawa, K. — UBC Sauder and NBER
Morlacco, M. — USC

July 2025

Abstract

This paper develops a theory of bargaining in firm-to-firm trade with two-sided market power. The framework accommodates flexible market structures, yielding analytical expressions for pair-specific markups and pass-through elasticities. In U.S. import data, we estimate strong importer bargaining power and an upward-sloping export supply curve, consistent with oligopsony power. Pass-through of the 2018 tariffs in firm-to-firm relationships is incomplete, in contrast to product-level studies, primarily due to exporter cost reductions driven by falling demand from dominant buyers. Our study highlights how bargaining and network rigidities shape price outcomes, with implications for markup dispersion and shock propagation in global value chains.

*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 Antrás, Costas Arkolakis, Ariel Burstein, Julia Cajal-Grossi, Thomas Chaney, Keith Head, Ali Hortacsu, Michael Irlacher, Oleg Itshkoki, Amit Khandelwal, 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

The recent resurgence of protectionist trade policies has reignited debate over how tariffs affect international prices. This question is particularly salient in today's global economy, where production is fragmented across global value chains and most trade occurs within ongoing relationships between importers and exporters. These firm-to-firm linkages shape the cross-border transmission of shocks (Boehm et al., 2019); however, their implications for price setting and tariff pass-through remain less well understood.

A growing body of literature highlights the limitations of standard trade models in this context. Rather than engaging in anonymous spot-market transactions, firms often trade within long-term relationships shaped by contract incompleteness, relationship-specific investments, and lock-in frictions (Antràs, 2015). In this environment, prices arise from bilateral negotiations between exporters and importers, both of whom may hold significant market power (Antràs and Staiger, 2012; Morlacco, 2019; Dhyne et al., 2022).

This paper develops a theory of bargaining in firm-to-firm trade that incorporates two-sided market power and network rigidities. The framework accommodates flexible upstream and downstream market structures and delivers closed-form expressions for markups and pass-through elasticities as functions of market shares and a small set of structural parameters. We estimate these parameters using U.S. import data, validate the model's cross-sectional predictions against observed pricing patterns, and use the estimated framework to quantify aggregate pass-through in firm-to-firm trade, highlighting how bargaining and network structure shape both the level and sources of pass-through in response to trade shocks.

Section 2 introduces the theory. We consider a fixed network of importer-exporter (or buyer-supplier) relationships, where each pair negotiates bilaterally over the price of an intermediate input. Given the price, the importer unilaterally chooses the traded quantity to minimize costs, making prices allocative. In the event of a disagreement, each party falls back on trade with its other existing partners. As a result, the network structure influences the value of each party's outside option, and in turn, their effective bargaining power.

Our framework offers a unified theory of exporter and importer market power in firm-to-firm trade. Exporter rents arise from three sources: *monopoly power* through product differentiation, *oligopoly power* due to supplier concentration, and *quasi-rents* from an upward-sloping residual supply curve. When exporters hold full bargaining power, the model nests standard trade settings in which prices reflect an *oligopoly markup* over marginal cost (Atkeson and Burstein, 2008; Dhyne et al., 2022). This markup increases with the exporter's share of the buyer's total input purchases—a measure we refer to as the exporter's *supplier share*.

Our first contribution is to characterize the role of importer market power in price-setting. This power arises when suppliers earn economic rents, creating scope for buyers to extract surplus by negotiating lower prices. It also reflects *oligopsony power*, as concentrated buyers can negotiate lower prices. When importers hold full bargaining power, prices can fall below marginal cost, generating a negative price–cost margin akin to an *oligopsony markdown*. This markdown declines with the importer's share of the supplier's total output–a measure we refer to as the importer's *buyer share*.[1] In general, the negotiated price reflects a convex combination of oligopoly and oligopsony forces, with weights determined by the firms' outside options and other exogenous determinants of bargaining power.

Our second contribution is to characterize how importer market power shapes the pass-through of tariff shocks to bilateral prices. We focus on short-run pass-through elasticities, derived as local perturbations around the static price equilibrium, while holding the network and general equilibrium conditions fixed. This isolates the intensive margin of price adjustment from longer-run responses driven by entry, reallocation, or macroeconomic feedback.

Importer market power affects pass-through through two distinct mechanisms: a *markup channel* and a *cost channel*. The markup channel captures how bilateral markups respond to price changes. When exporters hold full bargaining power, markups exhibit *strategic complementarities*, a well-known source of incomplete pass-through (Amiti et al., 2014, 2019a). With two-sided market power, the markup elasticity becomes a convex combination of strategic complementarities and *strategic substitutabilities*, as importers reduce markdowns in response to higher prices, amplifying pass-through.

The cost channel captures how tariffs affect the exporter's marginal cost through scale effects. While prior work links pass-through to cost adjustment under decreasing returns to scale (e.g., Burstein and Gopinath 2015), our framework introduces a new source of heterogeneity: the elasticity of residual supply increases with an importer's buyer share. As a result, larger buyers face more incomplete pass-through. Although both channels interact to shape price responses, the cost channel dominates when bargaining power is two-sided and returns to scale are decreasing. In this setting, markup adjustments are modest and offset by opposing strategic forces, while cost-side responses amplify the heterogeneity of pass-through.

Section 3 describes the data and presents preliminary empirical evidence. Our main data source is the U.S. Census Longitudinal Firm Trade Transaction Database (LFTTD), which reports prices and quantities for each import transaction involving a U.S. importer, along with identifiers for the foreign exporter and the 10-digit product code. These features allow

[1] Our definition of the importer's markdown differs from the standard price–marginal revenue wedge. See Section 2 for details.

us to construct bilateral market shares on both sides of the transaction, which are the key sufficient statistics in the model governing the distribution of markups and pass-through elasticities. Our main sample focuses on repeated, arm's-length relationships involving the exchange of intermediate inputs, excluding related-party links as flagged in the LFTTD.

Pricing and pass-through patterns in the data closely align with the model's predictions under two-sided market power. Controlling for supplier–product–time fixed effects to proxy for marginal costs, we find that prices rise with the exporter's supplier share and fall with the importer's buyer share, consistent with the presence of both oligopoly and oligopsony forces. In contrast, tariff pass-through declines with buyer share and is non-decreasing in supplier share, suggesting that price adjustments are primarily driven by the cost channel, with a more limited role for the markup channel, as expected under two-sided market power.

These empirical comovements reinforce the model's central mechanisms and are difficult to reconcile with conventional pricing frameworks that assume one-sided or purely competitive behavior. Nonetheless, reduced-form estimates may conflate the effects of competitor behavior, input price variation, or broader macroeconomic shocks, limiting their interpretability. A structural approach is therefore needed to isolate the underlying mechanisms.[2]

To do so, Section 4 develops a structural estimation strategy around two key parameters: the importer's bargaining power and the exporter's returns to scale, which determines the slope of the input supply curve. We estimate that U.S. importers hold bargaining power of approximately 0.8, roughly four times greater than that of their foreign suppliers, and that the returns to scale parameter is about 0.45, consistent with upward-sloping supply. Identification exploits cross-sectional price variation across U.S. importers within a supplier–product–year. As the estimating equations hold independently of the general equilibrium environment, the approach remains valid without explicitly modeling broader equilibrium feedback.

We assess the model's validity by evaluating its ability to replicate both the average and heterogeneous effects of the 2018 tariffs on prices and quantities. First, the model closely matches the average price response as well as its variation with buyer and supplier shares, consistent with the role of oligopoly, oligopsony, and cost-side adjustments. Second, applying the IV-based test proposed by Adão et al. (2023), we find that our model outperforms alternatives that either impose constant returns to scale or exclude bargaining. Third, while the model underestimates the magnitude of quantity changes, it successfully captures their direction, lending support to its core allocative mechanism. Together, these results validate the model's key assumptions and suggest that general equilibrium feedback or omitted strategic interactions are unlikely to materially bias its pass-through predictions.

[2]For a related discussion in a different context, see Berger et al. (2022).

In Section 5, we use the estimated model to assess how the 2018 Trump tariffs affected aggregate U.S. import prices through firm-to-firm relationships. Rather than providing a full accounting of tariff incidence, which lies beyond the scope of our framework, we focus on one key margin: price adjustment within ongoing matches, and its decomposition into changes in markups and marginal costs.

We find that aggregate pass-through in firm-to-firm trade is substantially incomplete, with elasticities ranging from 65 to 71%. The model sheds light on the mechanisms behind this muted response. The key insight is that most of the adjustment occurs on the cost side: exporters facing weaker demand from powerful buyers move down their marginal cost curves, absorbing a sizable share of the tariff burden. Markup changes occur but contribute little to the aggregate response due to offsetting effects of strategic complementarities and substitutabilities. This distinction is important. While prior work often attributes incomplete pass-through to variable markups and strategic pricing, our results highlight the dominant role of cost adjustment and supply-side forces.

Our results stand in contrast to studies documenting near-complete pass-through of the 2018 tariffs (e.g., Amiti et al., 2019b, 2020; Cavallo et al., 2020; Fajgelbaum et al., 2020; Flaaen et al., 2020). We reconcile this divergence on two grounds. First, our analysis focuses on intermediate inputs and arm's-length transactions, where pass-through may be lower due to input specificity, contractual frictions, or capacity constraints. Second, we isolate pricing within repeated matches between importers and exporters. By contrast, product-level unit values used in aggregate studies conflate within-match price changes with compositional shifts in trading partners, notably including one-off exchanges. Such one-off transactions are more likely to reflect marginal cost pricing and exhibit full pass-through.[3]

Related Literature Our paper contributes to the literature on pricing and shock transmission in firm-to-firm trade, with a particular focus on the role of importer market power. A growing body of empirical work highlights the influence of dominant buyers in shaping supplier outcomes. Bernard et al. (2019) and Bernard et al. (2022) document the pivotal role of large buyers in determining supplier performance and market access. Using French customs data, Fontaine et al. (2020) show substantial variation in unit values across importers transacting with the same exporter, consistent with buyer-specific pricing. Similar patterns are found by Huang et al. (2021) in France, Chile, and China. Among the few studies to model buyer power directly, Morlacco (2019) estimates substantial oligopsony power among French importers, while Atkin et al. (2024) show that bargaining between Argentinian importers

[3]See Cajal-Grossi et al. (2023) for evidence that markups are lower and prices more competitive in spot transactions than in relational matches.

and foreign suppliers affects the price impact of trade policy.

Formal models of pricing in firm-to-firm trade remain limited. Dhyne et al. (2022) develops a network-based model of oligopoly in which suppliers' markups increase with their relationship-specific market share. Our framework nests this model as a special case but extends it to incorporate bilateral bargaining and importer market power. Grossman et al. (2024) also study firm-to-firm pricing under Nash-in-Nash bargaining, emphasizing how supply-chain links adjust in response to trade shocks. In contrast, we take the trade network as fixed and focus on how bargaining and network rigidities interact to determine equilibrium prices and short-run pass-through.

We also contribute to the literature on the firm-level determinants of cost shock pass-through. A large body of work attributes incomplete pass-through to variable markups and strategic complementarities in price-setting (e.g., Amiti et al., 2014, 2019a).[4] We extend these insights to a bilateral oligopoly setting that reflects the structure of a large share of international trade. Our framework nests these existing mechanisms but allows for a richer characterization of both markup and cost channels, shaped by bargaining power and network position. Related work includes Gopinath and Itskhoki (2011) and Goldberg and Tille (2013), who model bargaining in firm-to-firm trade but abstract from importer market power and cost-side adjustment. Empirically, we find that these cost-based responses, rather than variable markups, are the primary driver of incomplete pass-through in these settings.

Finally, we contribute to the literature on shock propagation in production networks. A large body of work demonstrates that shocks can spread through input–output linkages and affect aggregate outcomes (Acemoglu et al., 2012; Di Giovanni et al., 2014; Grassi, 2018), often using natural experiments to trace transmission along supply chains (Barrot and Sauvagnat, 2016; Boehm et al., 2019; Carvalho et al., 2021). Recent work by Acemoglu and Tahbaz-Salehi (2025) highlights how bilateral bargaining can amplify supply chain disruptions through the price channel. We complement this perspective by modeling how idiosyncratic shocks to individual relationships affect prices under two-sided market power. While we abstract from general equilibrium feedback, our focus is on the within-network transmission of trade shocks through pricing, a key but underexplored margin in the propagation of shocks across production networks.

[4] See also Berman et al. (2012), Auer and Schoenle (2016), and Garetto (2016) for firm-level evidence consistent with this class of models.

2 Theoretical Framework

This section develops a bargaining theory of firm-to-firm trade with two-sided market power. The model links markups and pass-through elasticities to importer and exporter market shares, along with a small set of structural parameters that capture key features of market structure upstream and downstream. The model yields analytic structural equations and testable predictions, which we later bring to the data using U.S. customs records.

We impose the following assumptions to keep the model tractable. First, we treat the trade network as fixed. This assumption captures the lock-in effects in firm-to-firm trade and allows us to focus on the price-setting problem while abstracting from firms' decisions to form or sever links.[5] Second, we abstract from nominal rigidities, such as fixed-price contracts or currency denomination, as these are unlikely to materially affect tariff pass-through.[6] Third, we consider a static framework of single-product negotiations, despite our data reflecting repeated, multi-product interactions between firms.[7] We return to these features and discuss how we incorporate information on relationship duration in Section 3 and 4.

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 treated as given.

Exporters and Supply Exporter i produces q_i units of the unique input variety and sells them to all importers in $\mathcal{Z}_i$, where total output satisfies $q_i \equiv \sum_{j \in \mathcal{Z}_i} q_{ij}$, and q_{ij} denotes the quantity of the intermediate input purchased by importer j.

We assume that exporter i operates a short-run production technology with returns to scale parameter $\theta \in (0, 1]$. Marginal cost (or equivalently, the short-run supply curve) is given by:

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

where k_i captures exogenous factors such as the exporter productivity or foreign wages.

[5] This separation between extensive and intensive margin decisions is consistent with prior work showing that dynamic sourcing choices, while important for long-run outcomes, are not necessary to analyze the short-run price effects of trade shocks (Blaum et al., 2018).

[6] Evidence from Amiti et al. (2020) supports this view, showing similar short- and long-run pass-through rates of 2018 U.S. import tariffs. See Goldberg and Tille (2013) for a theory of importer-exporter bargaining over the transaction price and exchange rate exposure.

[7] Repeated relationships affect contract structure and trade volumes (Macchiavello and Morjaria, 2015), exchange rate pass-through (Heise, 2024), and the propagation of shocks (Martin et al., 2023).

This cost structure implies that the average cost of q_i units is θ-times the marginal cost. When $\theta < 1$, indicating decreasing returns to scale, the short-run average cost lies below the marginal cost, and the gap between them widens as q_i increases. In contrast, $\theta = 1$ indicates a constant returns technology, with marginal and average costs constant and equal to k_i.

Importers and Demand Importer j combines domestic and foreign inputs, denoted by q_j^d and q_j^f respectively, to produce q_j units of a final good, which is subsequently sold in a downstream market. The production technology of importer j is given by:

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

where φ_j is the importer's productivity, γ and $\varrho - \gamma$ are the output elasticities of foreign and domestic inputs, respectively, with ϱ governing the degree of returns to scale in production.[8]

The foreign input q_j^f is a constant elasticity of substitution (CES) composite of differentiated input varieties sourced from exporters in the set $\mathcal{Z}_j$. Each variety q_{ij} is weighted by a demand shifter ς_{ij}, and the elasticity of substitution across varieties is given by $\rho > 1$:

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

Downstream, each importer operates in a monopolistically competitive market, facing CES demand with constant elasticity $\nu > 1$ and an importer-specific demand shifter D_j.

2.2 Bargaining Protocol and Gains From Trade

Each $i - j$ vertical interaction proceeds in two stages. In the first stage, the importer chooses input quantity q_{ij} to minimize total input costs, taking the price p_{ij} as given:

$$\min_{q_{ij}} \; TC_j(q_{ij}; p_{ij}) = p_{ij} q_{ij} + \sum_{\ell \in \mathcal{Z}_j \setminus i} p_{\ell j} q_{\ell j} + p_j^d q_j^d.$$

The solution to this problem yields importer j's input demand:

$$q_{ij} = q_j^f\, \varsigma_{ij}^{\rho} \left(\frac{p_{ij}}{p_j^f}\right)^{-\rho} \quad \text{with} \quad q_j^f = \gamma\, c_j\, q_j \left(p_j^f\right)^{-1}, \tag{2.4}$$

[8]The domestic input q_j^d can be interpreted as a constant returns to scale aggregator of primary factors, such as labor and domestic intermediates. Due to the lack of data on domestic input purchases, we model foreign and domestic intermediates as distinct rather than symmetric or substitutable.

where $p_j^f = \left(\sum_i \varsigma_{ij}^{\rho} p_{ij}^{1-\rho}\right)^{\frac{1}{1-\rho}}$ is the foreign input price index.

In the second stage, the importer and exporter determine the negotiated price p_{ij} through Nash bargaining, taking the importer's demand in equation (2.4) as given:

$$\max_{p_{ij}} \left(\underbrace{\pi^i(p_{ij}; q_{ij}) - \tilde{\pi}^i_{(-j)}}_{GFT^i_{ij}} \right)^{1-\phi} \left(\underbrace{\pi^j(p_{ij}; q_{ij}) - \tilde{\pi}^j_{(-i)}}_{GFT^j_{ij}} \right)^{\phi}. \tag{2.5}$$

In equation (2.5), $\phi \in (0, 1)$ denotes the importer's bargaining leverage, and GFT^k_{ij} for $k \in \{i, j\}$ captures the gains from trade, defined as the difference between a firm's profits with and without the bilateral relationship. In case of disagreement, each party falls back on trade with its other existing partners. All bilateral matches generate strictly positive surplus, and trade occurs in equilibrium across all relationships.

To solve equation (2.5), we adopt the Nash-in-Nash bargaining solution concept, in which each bilateral negotiation takes as given the outcomes of all other matches in the network (Horn and Wolinsky, 1988; Collard-Wexler et al., 2019).[9] For this reason, we leave the dependence on prices and quantities in other links in the network implicit throughout the analysis.

This setup implies two key properties. First, prices are allocative: the negotiated price p_{ij} pins down the traded quantity q_{ij}. Second, the equilibrium allocation (p_{ij}, q_{ij}) satisfies the importer's demand in equation (2.4). We discuss the allocative implications of these assumptions, along with alternative bargaining protocols, in Section 2.5.

2.3 Equilibrium Prices and Allocations

The solution to (2.5) yields a bilateral price of the form $p_{ij} = \mu_{ij} \cdot c_i$, where c_i is exporter i's marginal cost, and μ_{ij} is a pair-specific markup.

We define two bilateral market shares that serve as key sufficient statistics in the model:

$$s_{ij} \equiv \frac{p_{ij} q_{ij}}{\sum_{k \in \mathcal{Z}_j} p_{kj} q_{kj}} \in [0, 1], \qquad x_{ij} \equiv \frac{q_{ij}}{\sum_{k \in \mathcal{Z}_i} q_{ik}} \in [0, 1],$$

[9] A common alternative is the sequential bargaining model of Stole and Zwiebel (1996), in which disagreement with one partner triggers renegotiation with others. This framework is suited to one-to-many settings such as labor markets, but is less appropriate for global supply chains, where firms engage in many-to-many relationships and isolated breakdowns rarely affect unrelated contracts. Moreover, it introduces additional complexity without yielding clear empirical gains in our context.

where s_{ij} denotes exporter i's *supplier share*, defined as its sales to importer j as a fraction of j's total foreign input expenditures, and x_{ij} denotes importer j's *buyer share*, equal to its purchases as a share of exporter i's total output.

The following proposition characterizes the equilibrium markup as a convex combination of two limiting pricing regimes:

Proposition 1. *The bilateral markup μ_{ij} can be expressed as*

$$\mu_{ij} = (1 - \omega_{ij}) \cdot \mu_{ij}^{\text{oligopoly}} + \omega_{ij} \cdot \mu_{ij}^{\text{oligopsony}}, \tag{2.6}$$

where the weight

$$\omega_{ij} = \frac{\frac{\phi}{1-\phi}\lambda_{ij}}{1 + \frac{\phi}{1-\phi}\lambda_{ij}} \in (0, 1) \tag{2.7}$$

represents the importer's effective bargaining power*, and $\lambda_{ij} \equiv -\frac{d \ln \pi_j}{d \ln(p_{ij}\, q_{ij})} \cdot \frac{\pi_j}{GFT_{ij}^j}$ captures the strength of the importer's outside option.*

Proof: See Appendix A.1.

Equation (2.6) shows that the equilibrium markup lies between two polar cases. When the exporter holds all the bargaining power ($\phi \to 0$), the markup converges to the oligopoly case $\mu_{ij}^{\text{oligopoly}}$. When the importer holds full bargaining power ($\phi \to 1$), it converges to the oligopsony case $\mu_{ij}^{\text{oligopsony}}$. For intermediate values, the outcome reflects the influence of both parties, with weight ω_{ij} determined by bargaining power and the importer's outside option (see Section 2.3.3).

We now examine these limiting cases in more detail and characterize how the bilateral markup depends on market shares and model primitives.

2.3.1 Oligopoly Markup

Under full exporter bargaining power ($\phi = 0$), the markup takes the standard form:

$$\mu_{ij}^{\text{oligopoly}} = \frac{\varepsilon_{ij}}{\varepsilon_{ij} - 1} \geq 1, \tag{2.8}$$

where ε_{ij} denotes the residual demand elasticity faced by exporter i, given by:

$$\varepsilon_{ij} = (1 - s_{ij}) \cdot \rho + s_{ij} \cdot \eta. \tag{2.9}$$

This elasticity is a weighted average of two components: the elasticity of substitution across foreign suppliers, ρ, and the elasticity η of the importer's foreign input bundle q_j^f with respect to its price index p_j^f:

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

Provided that $\rho > \eta$, a standard parameter condition, the exporter's markup (2.8) increases with its supplier share s_{ij}, reflecting greater oligopoly power.

This case corresponds to the firm-to-firm trade setting in Dhyne et al. (2022). It is also closely related to the oligopolistic competition model of Atkeson and Burstein (2008), but differs in three key respects. First, supplier shares are defined at the match level, rather than at the firm or industry level, reflecting our assumption of a fixed network. Second, whereas the outer nest elasticity η is typically treated as a fixed preference parameter, here it summarizes how downstream market structure–determined by both technology (γ, ϱ) and demand (ν)–shapes the bargaining environment. Third, we allow for decreasing returns to scale in production, in contrast to the constant returns assumption common in related models.

Figure 1 illustrates the equilibrium outcome for a representative $i-j$ match. The downward-sloping residual demand and marginal revenue curves (black and gray) intersect with the upward-sloping residual marginal and average cost curves (blue and red).

Panel (A) depicts the case with $\phi = 0$. The equilibrium quantity is determined by the intersection of marginal cost and marginal revenue, and the price lies on the residual demand curve. Exporter rents have two components: the *oligopoly rent* (red), which corresponds to the markup in equation (2.8), and the *quasi-rent* (purple), which arises from the gap between marginal and average cost under decreasing returns to scale.

2.3.2 Oligopsony Markdown

Under full importer bargaining power ($\phi = 1$), the bilateral markup simplifies to:

$$\mu_{ij}^{\text{oligopsony}} = \theta\cdot\frac{1-(1-x_{ij})^{\frac{1}{\theta}}}{x_{ij}} \leq 1, \tag{2.10}$$

which we refer to as an *oligopsony* markup, or *markdown*, since it lies weakly below one.[10]

[10] Although $\mu_{ij}^{\text{oligopsony}}$ denotes a price–cost ratio, we refer to it as a markdown because it falls below one. This differs from standard oligopsony models, where markdowns typically reflect a wedge between input prices and marginal revenue product. See Section 2.5 for further discussion.

FIGURE 1: Illustration of Equilibrium Allocations With Different Bargaining Power

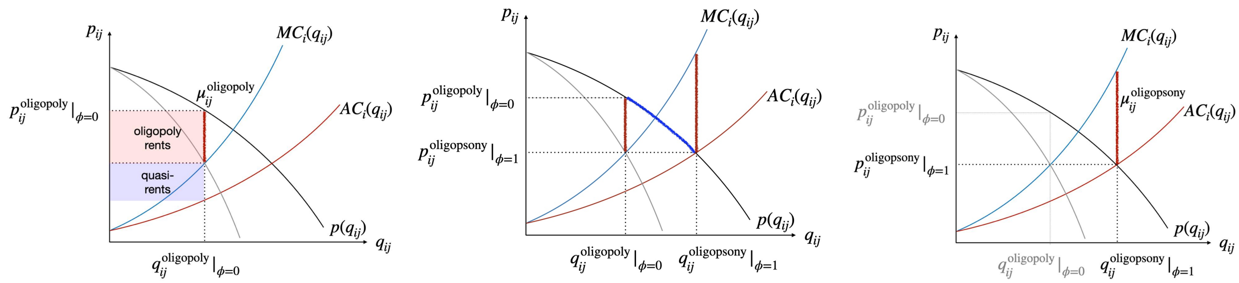

(A) $\phi = 0$ (B) $\phi \in (0, 1)$ (C) $\phi = 1$

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 $\theta < 1$.

Expression (2.10) characterizes the lowest price at which the exporter is willing to supply, with the importer extracting the exporter's entire surplus through bargaining.[11] The markdown declines with the buyer's share x_{ij}: larger buyers induce greater quasi-rents and can negotiate lower prices. As $x_{ij} \to 0$, the importer behaves atomistically, generates no quasi-rents, and pays a price equal to marginal cost.

This dependence of the price–cost ratio on x_{ij} reflects a form of *oligopsony power* that arises only under decreasing returns to scale. When $\theta = 1$, marginal and average costs coincide, quasi-rents vanish, and the importer can extract only oligopoly rents. In this case, the price cannot fall below marginal cost, and the markdown remains fixed at one.

Panel (C) of Figure 1 illustrates the equilibrium under $\phi = 1$. Compared to Panel (A), the price lies at the intersection of the downstream demand curve and the exporter's residual average cost curve, leaving no surplus to the exporter. The result is a lower price and higher quantity relative to the oligopoly case.

Clarification of Terminology The terms "buyer power" and "oligopsony power" are often used interchangeably in the literature, but rarely precisely defined (Noll, 2005). We distinguish between these concepts in our framework.

We use *buyer power* or *importer market 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 specific case when a higher buyer share x_{ij} enables the importer to negotiate lower prices, holding ϕ fixed. Importantly, oligopsony power requires both $x_{ij} > 0$ and $\theta < 1$, as it depends on the scale-dependent gap between marginal and average cost.

[11] Formally, the markdown is proportional to the percentage increase in quasi-rents attributable to the $i - j$ match. See equation (A.9) in Appendix A.1.

2.3.3 Two-Sided Market Power and Bargaining Weights

Panel (B) of Figure 1 shows the intermediate case with $\phi \in (0,1)$. The equilibrium price lies on the thick blue curve and reflects a weighted average of the outcomes in Panels (A) and (C), consistent with Proposition 1. The resulting price–marginal cost ratio can exceed or fall below one, depending on the relative bargaining power of the exporter and importer.

Determinants of ω_{ij}. The bargaining weight ω_{ij} in equation (2.6), defined in equation (2.7), governs the balance between oligopoly and oligopsony forces in price setting. It depends both on the exogenous bargaining power ϕ and an endogenous term λ_{ij}, defined as:

$$\lambda_{ij} \equiv \underbrace{-\frac{d\ln \pi_j}{d\ln(p_{ij}\, q_{ij})}}_{\text{Cost exposure } (\lambda_{ij}^C)} \cdot \underbrace{\frac{\pi_j}{GFT_{ij}^j}}_{\text{Network dependence } (\lambda_{ij}^N)} \geq 0.$$

To interpret λ_{ij}, we decompose it into two terms. The *cost exposure* component, $\lambda_{ij}^C = \frac{(\eta-1)s_{ij}}{\varepsilon_{ij}-1}$, captures how sensitive importer j's profits are to price changes from supplier i and increases with the supplier share s_{ij}. The *network dependence* term, $\lambda_{ij}^N = \left(1-(1-s_{ij})^{\frac{\eta-1}{\rho-1}}\right)^{-1}$, reflects how much of importer j's profits rely on the match with i, and declines with s_{ij} as reduced diversification weakens the importer's outside option.

Taken together, λ_{ij} captures both the importer's incentive to negotiate aggressively (through cost exposure) and the strength of its fallback option (through network dependence). It follows a hump-shaped pattern in s_{ij}, increasing at low values and declining at higher ones, and converges to 1 as $s_{ij} \to 0$ or $s_{ij} \to 1$.

This interpretation clarifies the meaning of ω_{ij} as a measure of the importer's *effective bargaining power*. While ϕ governs baseline influence in Nash bargaining, λ_{ij} endogenously adjusts it based on the structure of the trading relationship.

Testable Implications. Equation (2.6) writes the bilateral markup μ_{ij} as a convex combination of the oligopoly markup and the oligopsony markdown, with bargaining power as weight. Since these components scale with bilateral market shares, the model delivers clear predictions on how markups co-move with s_{ij} and x_{ij} under two-sided market power.

Proposition 2. *The bilateral markup μ_{ij} exhibits the following properties:*

1. **Markup and Exporter's Share.** *If μ_{ij} increases with s_{ij}, then $\phi < 1$.*

2. **Markup and Importer's Share.** *If μ_{ij} decreases with x_{ij}, then $\phi > 0$ and $\theta < 1$.*

Proof: See Appendix A.2.

Proposition 2 provides sufficient conditions for detecting two-sided market power. A positive relationship between μ_{ij} and supplier share s_{ij} implies $\phi < 1$, since exporters retain pricing power and $\mu_{ij}^{\text{oligopoly}}$ increases with s_{ij}. If $\phi = 1$, prices are fully disciplined by buyers, and no such pattern should emerge. While the dependence of ω_{ij} on s_{ij} may attenuate this effect, a positive slope remains inconsistent with full buyer power.

Conversely, a negative relationship between μ_{ij} and buyer share x_{ij} signals oligopsony power. This requires both positive bargaining leverage for importers ($\phi > 0$) and decreasing returns to scale ($\theta < 1$). When both conditions are met, markups reflect the joint influence of importer and exporter market power.

2.4 Equilibrium: Tariff Pass-Through Elasticities

Our pricing framework provides a basis for analyzing the short-run impact of an unanticipated tariff imposed on imports from country c. Let T_c denote the gross tariff rate. The (log) price that exporter i from country c charges importer j is given by:

$$\ln p_{ij} = \ln \mu_{ij} + \ln c_i + \ln T_c.$$

While a tariff on country c may, in principle, affect all trade relationships involving exporters from c, we focus on the short-run, direct effect at the relationship level. We interpret the tariff shock as a small, unanticipated perturbation around the initial equilibrium. Accordingly, we treat T_c as a pair-specific shock and hold constant prices in all other relevant matches such that $dp_{rj} = 0$ 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:*

$$\Phi_{ij} \equiv \frac{d \ln p_{ij}}{d \ln T_c} = \frac{1}{1 + \Gamma_{ij} + \Lambda_{ij}}, \tag{2.11}$$

where:

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

denote the partial elasticities of the equilibrium markup μ_{ij} and the exporter's marginal cost c_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 pricing responses, and a cost channel, capturing how the exporter's marginal cost adjusts in response to price changes. In models with CES demand, monopolistic competition, and constant marginal costs, both Γ_{ij} and Λ_{ij} are zero, yielding full pass-through ($\Phi_{ij} = 1$). More generally, Φ_{ij} may lie above, below, or at one depending on the magnitudes of Γ_{ij} and Λ_{ij}.

2.4.1 Markup Elasticity

The markup elasticity Γ_{ij} measures how the equilibrium markup μ_{ij} responds to changes in the bilateral price p_{ij}. It is given by:

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

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

The term in square brackets captures the direct elasticity of the markup, expressed as a convex combination of the oligopoly and oligopsony components. The second term reflects how the bargaining weight ω_{ij} itself responds to price changes. We discuss each component in turn.

Oligopoly Markup Elasticity The oligopoly markup elasticity dominates when $\phi \to 0$, in which case $\Gamma_{ij} \to \Gamma_{ij}^{\text{oligopoly}}$. It is given by

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

This elasticity reflects the standard logic of *strategic complementarities* in price-setting: when a tariff raises the bilateral price p_{ij}, exporter i reduces its markup to limit trade diversion, resulting in incomplete pass-through (see Amiti et al., 2014; Auer and Schoenle, 2016; Garetto, 2016; Amiti et al., 2019a). The response is strongest at intermediate supplier shares, giving rise to a U-shaped relationship between $\Gamma_{ij}^{\text{oligopoly}}$ and s_{ij}.[12]

[12] Amiti et al. (2014) shows that, to a first-order approximation, pass-through decreases with the exporter's share through strategic complementarities. However, we cannot rely on the same approximation, given our focus on bilateral markets, where both very low and very high market shares are observed in the data.

Oligopsony Markdown Elasticity The oligopsony markdown elasticity dominates when $\phi \to 1$, in which case $\Gamma_{ij} \to \Gamma_{ij}^{\text{oligopsony}}$. It is given by

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

This elasticity captures the logic of *strategic substitutabilities* among importers.[13] As the price p_{ij} rises, importer j reduces demand, shrinking the exporter's quasi-rents. This weakens the importer's bargaining position, reducing the markdown and amplifying the price response, potentially resulting in more-than-complete pass-through.

The elasticity $\Gamma_{ij}^{\text{oligopsony}}$ depends on both the importer's buyer share x_{ij} and the exporter's supplier share s_{ij}. It is U-shaped in x_{ij}: the elasticity vanishes when $x_{ij} \to 0$ (atomistic buyer) or $x_{ij} \to 1$ (monopsonist), and reaches its peak at intermediate values. It also declines with s_{ij}, since a higher supplier share reduces the demand elasticity ε_{ij}, limiting changes in x_{ij} and weakening the markdown response.

The Role of the Endogenous Bargaining Weight The elasticity $\Gamma_{ij}^{\omega} \equiv \frac{d \ln \omega_{ij}}{d \ln p_{ij}}$ captures how the bargaining weight ω_{ij} responds to price changes. Its sign and magnitude depend on the exporter's supplier share s_{ij} and are derived in Appendix A.3. For empirically relevant values of s_{ij} and parameter ranges, this elasticity is typically small. In particular, as s_{ij} approaches zero or one, where ω_{ij} converges to the importer's exogenous bargaining strength ϕ, the elasticity Γ_{ij}^{ω} goes to zero.

2.4.2 Cost Elasticity

The cost elasticity Λ_{ij} measures how exporter i's marginal cost c_i responds to changes in the bilateral price p_{ij}, via adjustments in traded quantity. While prior work emphasizes the role of the cost channel in shaping pass-through (e.g., Burstein and Gopinath, 2015; Amiti et al., 2019a), our contribution is to show that this elasticity varies systematically across matches through both demand and supply channels.

Formally:

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

[13] Strategic substitutabilities arise because a decline in demand by other buyers lowers marginal costs and prices, encouraging importer j to expand purchases.

FIGURE 2: Pass-Through Elasticity and Bilateral Market Shares

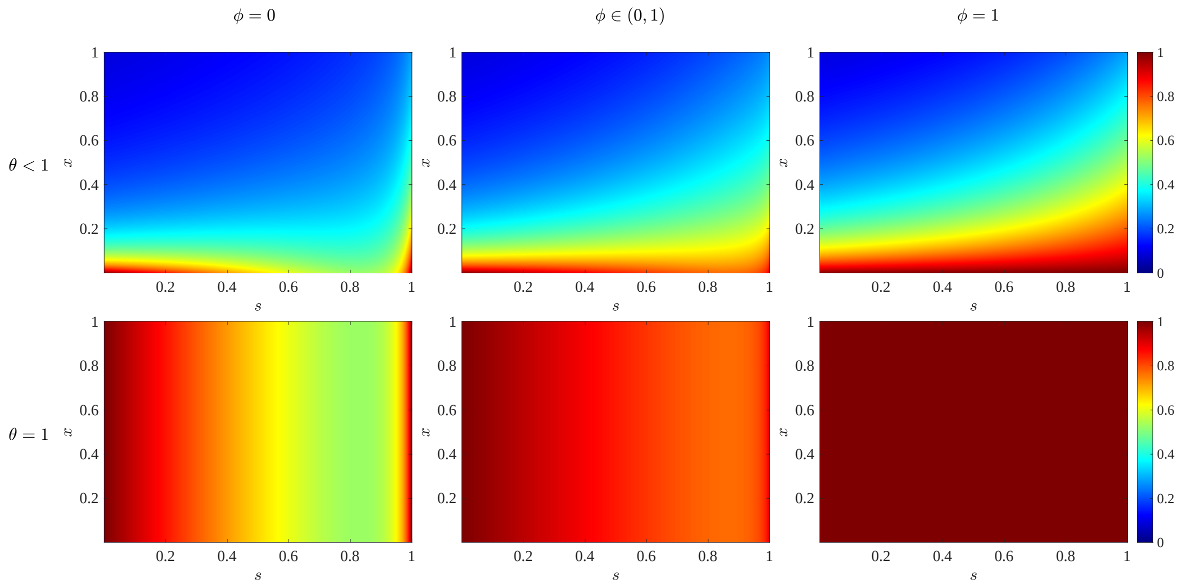

Notes: The figure presents heatmaps of the pass-through elasticity Φ_{ij} across combinations of s_{ij} (x-axis) and x_{ij} (y-axis), under alternative assumptions about ϕ and θ. We set $\theta = 0.5$ in the decreasing returns case and $\phi = 0.5$ in the intermediate bargaining case. Other parameters are fixed at $\gamma = 0.5$, $\varrho = 1$, $\nu = 4$, and $\rho = 10$.

The elasticity declines with the exporter's supplier share s_{ij}, which reduces the residual demand elasticity ε_{ij}, and increases with the importer's buyer share x_{ij}, which amplifies the residual supply elasticity $\frac{d \ln c_i}{d \ln q_{ij}}$. Importantly, this mechanism operates independently of bargaining power (ϕ), and instead reflects the concentrated nature of the trade network.

2.4.3 Pass-Through Elasticity and Bilateral Market Shares

The interaction between markup and cost channels generates rich heterogeneity in pass-through elasticities across matches, shaped by bargaining power (ϕ), returns to scale (θ), and bilateral market shares (s_{ij} and x_{ij}). Figure 2 illustrates these patterns using heatmaps of the pass-through elasticity Φ_{ij} as a function of supplier share (s_{ij}) and buyer share (x_{ij}), under three bargaining regimes—$\phi \to 0$, $\phi \in (0,1)$, and $\phi \to 1$—with decreasing returns to scale ($\theta < 1$) in the top row and constant returns to scale ($\theta = 1$) in the bottom row.

Several key insights emerge. First, pass-through Φ_{ij} increases with importer bargaining power ϕ. As ϕ rises, the markup elasticity shifts toward the oligopsony markdown elasticity, increasing pass-through through strategic substitutabilities. This is most visible in the shift from the left to the right columns within each row.

Second, when $\theta < 1$, pass-through is incomplete across most of the (s_{ij}, x_{ij}) space. This is because the cost elasticity Λ_{ij}, which is always non-negative, dominates the overall pass-through elasticity ($\Gamma_{ij} + \Lambda_{ij}$) when Γ_{ij} is negative due to strong buyer power. Moreover, pass-through values under $\phi \in (0, 1)$ and $\phi = 1$ are nearly indistinguishable, indicating that pass-through is relatively insensitive to the exact level of ϕ in this regime. This reflects two forces: (i) the cost channel is strong when returns to scale are decreasing and importers' bargaining power is high, and (ii) the markup elasticity Γ_{ij} is either muted by offsetting effects of strategic complementarities and substitutabilities or generally low in values.[14]

Third, pass-through Φ_{ij} depends on x_{ij} only when $\theta < 1$, in which case it declines with x_{ij}. This yields a sharp empirical prediction: a negative relationship between pass-through and buyer share x_{ij} implies decreasing returns to scale, regardless of the level of bargaining power. This result is formalized below.

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

Proof: See Appendix A.4.

By contrast, the relationship between Φ_{ij} and the supplier share s_{ij} is less clear-cut. As shown in Figure 2, pass-through declines with s_{ij} primarily when ϕ is low and θ is near one, a setting in which the markup channel dominates and cost-based adjustments are limited. In this region, strategic complementarities give rise to a U-shaped relationship between Φ_{ij} and s_{ij}. As either ϕ or x_{ij} increases, the cost channel becomes more prominent, and pass-through tends to rise with s_{ij}. While we do not formally characterize the conditions under which this reversal occurs, we show below that ϕ is weakly increasing in s_{ij}. This monotonicity supports interpreting the empirical evidence as inconsistent with the low-ϕ, high-θ case.

2.5 Discussion

This section reviews key modeling assumptions and discusses potential extensions.

2.5.1 Bargaining Protocol and Quantities

The baseline model assumes *demand-determined quantities*, whereby the importer chooses input quantities to minimize total cost, taking the price as given. This assumption provides analytical tractability, yields closed-form solutions for markups and pass-through, and nests standard models of international trade. An important implication is that prices are allocative, a feature that aligns with our empirical findings and with firm-level evidence in related settings (Gopinath and Itskhoki, 2011).

We consider two alternative bargaining protocols as benchmarks for understanding the role

[14] For more details, see Appendix A.4.

of quantity determination. The first is *efficient bargaining*, discussed more formally in Appendix B.1. This case corresponds to the vertically integrated benchmark, where importer and exporter jointly negotiate over both price and quantity to maximize total surplus. While theoretically appealing, this setup implies that prices are non-allocative transfers, inconsistent with the empirical evidence. Moreover, the vertically integrated case may be unrealistic in the context of arm's-length firm-to-firm trade, where limited commitment and contracting frictions are prevalent (Antràs, 2020).

The second alternative, detailed in Appendix B.2, is *supply-driven* bargaining: the exporter sets quantity for a given price, and price is then negotiated based on the resulting supply curve. As $\phi \to 1$, this nests the classic monopsony benchmark commonly used in labor markets (e.g., Berger et al., 2022). The key distinction between this setup and our baseline lies in their welfare implications: in the supply-driven case, buyer power lowers both prices and quantities, while supplier power helps restore efficiency. In contrast, under demand-driven bargaining, buyer power mitigates upstream distortions and improves efficiency.[15]

While both alternatives offer useful benchmarks, they are less suited to the goals of this paper. The supply-driven model, in particular, does not provide closed-form expressions for key variables, which makes it challenging to generate clear predictions or link the model to data compared to our baseline setup. Still, it remains a valuable direction for future work, especially for studying welfare effects when buyer power leads to inefficiencies.

2.5.2 Outside Options

In our model, each firm's outside option reflects payoffs from trade with all other existing partners, excluding the focal match. This assumption allows us to express markups and pass-through in terms of market shares and a small set of parameters, facilitating structural estimation and counterfactual analysis.

While analytically convenient, the assumption may appear restrictive if disagreement leads a firm to form or sever other relationships. Appendix B.3 explores a more flexible setup in which disagreement affects the importer's cost and the exporter's revenue non-parametrically. Although more realistic, this extension introduces an identification problem: the parameters ϕ and θ can no longer be separately identified from the outside option, which limits their interpretability and empirical tractability.

[15] See Avignon et al. (2024) and Demirer and Rubens (2025) for recent discussions of how quantity-setting assumptions affect welfare outcomes.

2.5.3 General Equilibrium Forces

We conclude the theoretical section by noting that all results are derived under a partial equilibrium approach, holding fixed general equilibrium variables such as wages, demand shifters, and competitor behavior. This approach simplifies the analysis but also raises questions about the role of general equilibrium adjustments in the results.

General equilibrium considerations are critical for interpreting pass-through elasticities. The structural pass-through elasticity derived in Section 2.4 is a *direct* or *partial* elasticity, holding constant aggregate variables such as wages, demand conditions, and, importantly, competitors' prices and sourcing decisions. In reality, tariff shocks may also affect variables such as foreign wages or domestic export prices. If these effects are not fully observed or controlled for, reduced-form estimates may conflate the direct impact with general equilibrium responses, making it harder to align empirical pass-through coefficients with model-based elasticities (Burstein and Gopinath, 2015).[16] Therefore, our goal is not to recover structural elasticities from reduced-form regressions, which would be inappropriate in this context (Berger et al., 2022).

Instead, we pursue a structural approach, which helps mitigate similar concerns. Our approach has three advantages. First, the theoretical relationship between bilateral shares, markups, and pass-through holds parametrically regardless of the specific general equilibrium environment. While the parameter values depend on how aggregate variables such as tariffs, demand shifters, and wages co-move in equilibrium, we do not need to model their relationships explicitly. Instead, identification is based on the cross-sectional variation of prices across importers within exporter–product–year cells, without making assumptions on the general equilibrium environment.

Second, we can also isolate and test the cross-sectional predictions of the model independently of general equilibrium forces. In particular, even if aggregate variables shift over time, the model predicts specific patterns between firm-to-firm shares, markups, and pass-through at a given point in time. These predictions can be assessed empirically using within-period variation across matches, which we exploit in the next sections.

Third, we can evaluate the performance of the estimated model by comparing its ability to replicate observed price changes in response to tariff shocks, thereby gauging the significance of general equilibrium forces. As shown in the empirical analysis, the estimated model fits the

[16] In Appendix B.4, we extend equation (2.11) to account for indirect effects, such as how a shock to exporter i influences other prices and quantities, which may in turn affect p_{ij}. While we do not model full general equilibrium dynamics, this extension illustrates how spillovers across relationships may lead to reduced-form estimates diverging from structural ones.

observed price changes well. This suggests that the short-run effects of tariffs on prices can be understood primarily through the lens of partial equilibrium mechanisms, and supports the usefulness of our framework for studying firm-to-firm pricing and pass-through.

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 evidence testing the model's predictions.

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, 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.[17]

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

[17]The MID combines the country code, (elements from) the firm name, city, and address (Kamal and Monarch, 2018). Because the algorithm is not standardized, it may generate inconsistent identifiers due to misspellings or minor location changes, leading to one firm having multiple MIDs or several firms sharing one. Following Kamal and Monarch (2018), we construct a robustness version that truncates location fields to improve consistency. Our baseline uses the full MID, and results are robust to this alternative.

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:

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

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 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 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 $x_{ijh} \equiv \frac{q_{ijh}}{\sum_{k \in \mathcal{Z}_i^h} q_{ihk}}$, 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.[18]

3.3 Sample Construction and Summary Statistics

We apply a series of restrictions to the LFTTD to align the empirical sample with the model's focus on decentralized bargaining over intermediate inputs. Full details are provided in Appendix C.2. The selection criteria are designed to ensure that we observe relationship-

[18] To address the possibility that the importer's buyer share x_{ij} may be overstated due to unobserved sales to other destinations, we replicate the analysis using only exporters from Canada and Mexico. These countries direct the majority of their exports to the U.S. (71% for Canada and 73% for Mexico in 2019), making the assumption of destination-specific production less restrictive. Reassuringly, our estimates remain stable in this subsample. Full results are available upon request.

level price changes, exclude related-party transactions, and maintain sufficient variation in buyer–supplier matches for identification purposes. Appendix Table C.1 summarizes how the sample evolves with each restriction across key dimensions.

We begin by restricting the sample to importer–exporter–product triples observed in two consecutive years, allowing us to compute relationship-level price changes, which are central to our analysis. While this requirement eliminates over half of the raw matches, it retains a substantial share of trade, covering roughly 88% of import value between 2001–2016 and 80% between 2017–2018, highlighting the importance of repeated relationships.

Next, we restrict the sample to capital and intermediate inputs by excluding products classified as consumption goods under the Broad Economic Categories (BEC) system. This step reduces the number of suppliers and relationships, as shown in Table C.1.[19] To further mitigate measurement error, we apply three filters: (i) exclude transactions involving energy goods, (ii) drop observations with unit values outside the 1st–99th percentile within product, and (iii) remove transactions with absolute log price changes above four.

We then exclude related-party transactions, which are less likely to reflect decentralized bargaining and may involve internal pricing strategies such as transfer pricing.[20] In our baseline definition, a buyer–supplier pair is considered related if ORBIS identifies a shared corporate parent.[21] For robustness, we consider two alternative definitions: one based solely on the LFTTD's related-party flag, and another that combines this flag with ORBIS data identifying the U.S. importer as a multinational (domestic or foreign-owned).

Lastly, we impose restrictions based on our identification strategy. Since the latter relies on observing the variation in prices across U.S. buyers for the same supplier-product-year combination, we restrict our sample to supplier–product pairs in which the supplier transacts with at least two U.S. buyers in consecutive years. Table C.1 reports that after these restrictions, the sample accounts for approximately $160 billion in import value and 250 thousand buyer-supplier-product-year combinations between 2017-2018.

[19] As a robustness check, we also consider a broader sample that includes consumption goods. Appendix Table C.2 reports the composition of this sample and Table C.3 reports the corresponding summary statistics. These statistics are broadly similar to the baseline, indicating that including consumption goods does not significantly alter the data composition.

[20] Bernard et al. (2006) document that prices in related-party trade differ systematically from arm's-length transactions, with lower average prices and distinct pass-through behavior.

[21] We retain all observations not flagged as related in either ORBIS or the LFTTD. This approach preserves sample size while ensuring that limited ORBIS coverage does not unduly constrain the selection. See Appendix C.1 for details.

Summary Statistics Table 1 reports summary statistics for our final sample. Panel A shows that the concentration of importers and exporters is substantial. 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%. The two shares are highly dispersed and largely uncorrelated, with a correlation of 0.04.

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

Variable	Mean	Std. Dev.	P25	Median	P75
Panel A: Characteristics of Trade Relationships					
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	—	—	—	—
Panel B: Prices					
$\log p$ (pre-duty)	3.50	2.80	1.40	3.10	5.40
$\log p$ (pre-duty, excl. charges)	3.40	2.80	1.30	3.00	5.40
$\log p^{\text{duty}}$ (post-duty)	3.50	2.80	1.40	3.10	5.40

Notes: This table reports summary statistics for the estimation sample used in the empirical analysis. The data span 2001–2016 and include importer–supplier–product matches observed in two consecutive years. The sample excludes consumption goods (based on BEC), energy products, statistical outliers, and related-party trade, and is restricted to suppliers trading with at least two U.S. buyers in two consecutive years. This sample corresponds to the "+ Supplier Multi-Buyer" row in Panel B of Table C.1. Columns report the mean, standard deviation, and selected quantiles (25th, 50th, 75th percentile) for each variable. Prices in Panel B are log unit values (FOB value over quantity), with variants including charges or duties. s_{ijh} denotes exporter i's share in buyer j's imports of product h; x_{ijh} denotes buyer j's share in exporter i's U.S. exports of the same product. Relationship length and tenure are in years; concentration is measured at the HS10–year level. Counts of buyers, suppliers, and origin countries are per product per firm. Statistics are based on confidential LFTTD data and rounded to four significant digits per U.S. Census Bureau Disclosure Guidelines. Source: FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

Long-term relationships are a hallmark of intermediate input trade (Antràs and Chor, 2013; Monarch, 2022). In our data, importer–exporter pairs trade the same product for an average of four years and remain connected across all products for nearly five years. Pairs transact frequently (median: 16 times per product), often spanning multiple products (mean: 3.8), and typically involve durable links, with average tenures ranging from 7 to 10 years.

Panel B of Table 1 reports descriptive statistics for three bilateral (log) price measures based on FOB unit values: baseline pre-duty prices, pre-duty prices excluding insurance and other ancillary charges, and post-duty prices. All three exhibit wide dispersion across importer–exporter–product matches, with interquartile ranges exceeding four log points.

To assess the sources of this variation, we perform a variance decomposition exercise, as described in Appendix D.1. Table D.1 shows that product-year fixed effects explain about 50% of the total variance, while match-specific residuals account for 4%. Crucially, when we focus only on variation within supplier–product–year cells, 77% of the remaining dispersion is explained by match-specific factors. This pattern holds across price definitions, underscoring the importance of relationship-specific forces in pricing.[22]

3.4 Test of Model Predictions

We now examine how the comovements between markups, pass-through, and bilateral market shares align with the model's predictions in Propositions 2 and 4. While not a formal test, this analysis provides supporting evidence for the mechanisms emphasized in the theory. Section 4 then develops a structural approach to quantify these forces more directly.

3.4.1 Test of Proposition 2: Markups and Bilateral Market Shares

We begin by testing Proposition 2, which predicts that with two-sided market power ($\phi \in (0,1)$ and $\theta < 1$), bilateral markups increase with the exporter's supplier share (s_{ijht}) and decrease with the importer's buyer share (x_{ijht}). Since markups are not observed, we use log prices, equal to log markups plus log marginal costs, and include supplier–product–time fixed effects to absorb cost variation and isolate the markup component.

We estimate the following specification:

$$\ln p_{ijht} = \alpha_s \, s_{ijht} + \alpha_x \, x_{ijht} + \mathbf{X}_{ijht}\boldsymbol{\gamma} + \mathbf{FE} + \upsilon_{ijht}, \tag{3.1}$$

where the coefficients of interest are α_s and α_x, which we expect to be positive and negative, respectively.

[22]See Fontaine et al. (2020) for related evidence in French data.

To address endogeneity concerns, we construct leave-one-out instruments that isolate variation in market structure plausibly exogenous to the pricing decision of a given buyer–supplier pair. Specifically, we use the average supplier share (excluding i) among other buyers of exporter i (excluding j) to instrument for s_{ijh}. Similarly, to instrument for x_{ijh}, we use the average buyer share (excluding j) among other suppliers to importer j (excluding i).

Table 2 shows the results. Columns (1)–(2) include exporter (FE_i), importer (FE_j), and product–year (FE_{ht}) fixed effects. Column (3)–(4) replaces exporter fixed effects with exporter–product–year (FE_{iht}) to account for unobserved marginal costs. Columns (5) and (6) further replace importer fixed effects with importer–product–year (FE_{jht}) to capture buyer-specific demand shocks. All regressions control for the relationship duration, measured as years since the first shipment of product h between i and j.

We estimate $\alpha_s > 0$ and $\alpha_x < 0$, with both coefficients statistically and economically significant. A positive α_s indicates oligopoly power, requiring exporter bargaining power ($\phi < 1$), while a negative α_x reflects oligopsony power, requiring importer bargaining power ($\phi > 0$) and upward-sloping supply ($\theta < 1$). These findings are consistent with the model's core assumption of two-sided market power, i.e., $\phi \in (0, 1)$ and $\theta < 1$.

3.4.2 Test of Proposition 4: Pass-Through and Bilateral Market Shares

We next test Proposition 4, which links tariff pass-through to the importer's buyer share. Under decreasing returns ($\theta < 1$), the model predicts that pass-through declines with the buyer share via the cost channel, providing a direct test for $\theta < 1$.

For this analysis, we focus on 2017–2018, when U.S. imports experienced sharp and unanticipated tariff increases under the Trump administration. We estimate the following regression specification:

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

where we use the change in the duty-exclusive price as the dependent variable, defined as $\Delta \ln p_{ijht} = \Delta \ln p^{\text{duty}}_{ijht} - \Delta \ln(1 + \tau^{\text{app}}_{cht})$, where τ^{app}_{cht} is the applied ad-valorem tariff. This transformation isolates price changes net of applied duties, mitigating measurement error.

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

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}	0.183	0.174	0.251	0.169	0.269	0.681
	(0.0033)	(0.0296)	(0.0036)	(0.029)	(0.0065)	(0.0147)
x_{ijht}	-0.537	-0.077	-0.594	-0.186	-0.533	-0.777
	(0.0030)	(0.0175)	(0.0028)	(0.0249)	(0.0061)	(0.0112)
$\text{FE}_i + \text{FE}_j + \text{FE}_{ht}$	Yes	Yes	No	No	No	No
$\text{FE}_{iht} + \text{FE}_j$	No	No	Yes	Yes	No	No
$\text{FE}_{iht} + \text{FE}_{jht}$	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	0.957	0.010	0.976	0.032	0.991	0.032
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 FOB 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}), flexibly controlling for sourcing and pricing patterns. 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), which assess instrument strength for each endogenous regressor conditional on the others, addressing limitations of standard first-stage tests in multi-equation IV settings. 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).

Table 3: Pass-Through and Relationship Heterogeneity

Dependent variable:	$\Delta \ln p_{ijht}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1+\tau_{cht})$	-0.151	-0.188	-0.171	-0.045	-0.066	-0.123
	(0.093)	(0.105)	(0.095)	(0.099)	(0.093)	(0.104)
$\Delta \ln(1+\tau_{cht}) \cdot \ln \text{longevity}_{ijht}$		0.026				0.042
		(0.019)				(0.018)
$\Delta \ln(1+\tau_{cht}) \cdot s_{ijht-1}$			0.050		0.054	0.053
			(0.075)		(0.072)	(0.071)
$\Delta \ln(1+\tau_{cht}) \cdot x_{ijht-1}$				-0.403	-0.403	-0.411
				(0.113)	(0.113)	(0.114)
$\text{FE}_{ht} + \text{FE}_{cs}$	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 duty-exclusive 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 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. See Table D.11 for results using an alternative definition of arm's-length trade based on LFTTD. Source: FSRDC Project Number 2109 (CBDRB-FY25-P2109-R12520).

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, adding 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 1.5% , corresponding to 85% pass-through rate. Column (2) adds an interaction with relationship age to account for the role of match longevity, which has been shown to influence price adjustments. The results confirm that pass-through rises with relationship length, consistent with the results in Heise (2024).

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

To assess nonlinearities, we interact tariff changes with quartiles of lagged 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 fixed effects.

Together, these results suggest that pass-through is largely unresponsive to supplier concentration but declines strongly with buyer concentration. This pattern aligns with the model's predictions under two-sided market power: the strong buyer share gradient provides direct evidence of decreasing returns, as formalized in Proposition 4, while the weak supplier share gradient reflects the dominance of cost-channel, which suggests a high-ϕ, low-θ environment.

4 Structural Estimation

The patterns documented above are consistent with the model's predictions under two-sided market power. However, as discussed, they do not warrant direct inference on the structural parameters. We now turn to a structural estimation approach to quantify the role of two-sided market power in shaping international prices.

Our estimation targets two key parameters: the importer's bargaining power, ϕ, and the returns to scale parameter, θ, which governs the elasticity of foreign export supply. The remaining parameters are taken from the literature or directly measured from the data.

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 adopt similar values to match observed U.S. markups. The downstream demand elasticity faced by importers is set to $\nu = 4$, based on the estimates in Broda and Weinstein (2006).[23] The elasticity of importer j's marginal cost with respect to the foreign input price index is set to $\gamma = 0.5$, following Eldridge and Powers (2018), who document the share of imported inputs in total material costs for U.S. manufacturers. Input cost shares α_{jht} are directly measured from the data.

Finally, we set the importer's returns to scale parameter to $\varrho = 1$ in the baseline specification. As discussed in Section 2.3.1, the parameters ν, γ, and ϱ jointly determine the elasticity η of importer j's foreign input bundle q_j^f with respect to its price index p_j^f. Since ϱ enters the

[23] Appendix E.1 provides further discussion.

model only through η, fixing it to one does not restrict generality. We verify that our results are robust to alternative values of η.

4.1 Identification and Estimation of the Parameters θ and ϕ

Let Ω_{ijt} denote the information set available to a given $i-j$ pair during negotiations. This includes observed market shares (s_{ijht}, x_{ijht}, α_{jht}) and calibrated parameters (ν, γ, ρ, ϱ). As shown in equation (2.6), the bilateral markup depends only on the model primitives (ϕ, θ) conditional on Ω_{ijt}, i.e., $\mu_{ij} = \mu\left(\phi, \theta \mid \Omega_{ijt}\right)$. The log price of product h transacted between exporter i and importer j in year t can thus be written as:

$$\ln p_{ijht} = \ln \mu\left(\phi, \theta \mid \Omega_{ijt}\right) + \ln c_{iht},$$

where $\ln c_{iht}$ denotes the exporter's marginal cost. From equation (2.1), this is given by:

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

In the theoretical model, the term k_{iht} captured exporter-level cost shifters. In the empirical implementation, we generalize this term to allow for match-specific cost components, such as relationship-specific know-how or specialization, by letting it vary flexibly at the match-year level, k_{ijht}, thereby absorbing all (i, j, h, t)-specific variation. In contrast, the term $\frac{1-\theta}{\theta} \ln q_{iht}$ is constant across importer matches for a given supplier–product–year. Identification of (ϕ, θ) thus relies on cross-sectional variation in bilateral market shares, as discussed next.

Consider an exporter i matched with two importers, j and ℓ. Conditional on the joint information set $\boldsymbol{\Omega}_{ij\ell t} \equiv (\Omega_{ijt}, \Omega_{i\ell t})$, we assume that the unobserved component of marginal cost is mean-independent of the buyer identity: $\Delta k_{ij\ell ht} \equiv \mathbb{E}_k\left[k_{ijht} - k_{i\ell ht} \mid \boldsymbol{\Omega}_{ij\ell t}\right] = 0.$[24] Taking log price differences across buyers j and ℓ served by the same exporter i in year t yields the following moment condition:

$$g(\phi, \theta \mid \boldsymbol{\Omega}_{ij\ell t}) \equiv \mathbb{E}_k\Big[\ln p_{ijht} - \ln p_{i\ell ht} - \Big(\ln \mu(\phi, \theta \mid \Omega_{ijt}) - \ln \mu(\phi, \theta \mid \Omega_{i\ell t})\Big)\Big|\boldsymbol{\Omega}_{ij\ell t}\Big] = 0. \quad (4.1)$$

Identification requires that equation (4.1) does not hold for two pairs (ϕ, θ) such that $(\phi^A, \theta^A) \neq (\phi^B, \theta^B)$. Since the oligopoly markup is independent of the returns to scale

[24]Omitted variables may induce $\Delta k_{ij\ell ht} \neq 0$, raising endogeneity concerns. One possible issue is endogenous network formation: unobserved factors may simultaneously affect both match formation and pricing, biasing estimates of $\hat{\phi}$ and $\hat{\theta}$. We address this by differencing across buyers of the same exporter, which removes exporter-level shocks common to all matches. To address remaining concerns, we further implement an instrumental variable strategy.

parameter θ, identification of θ requires that the oligopsony channel plays a role in price determination, i.e., $\phi > 0$, a condition supported by the reduced-form evidence in Tables 2 and 3.

We therefore focus on the empirically relevant case of bilateral bargaining power $\phi \in (0, 1)$. The markup function $\mu(\phi, \theta \mid \Omega)$ is strictly monotonic in both parameters and thus invertible in each. It follows that the moment condition in equation (4.1) is also invertible in ϕ and θ. Identification then relies on observing multiple importer–exporter pairs in the same year, or multiple matches for a given exporter over time, under the assumption that bargaining weights remain constant across matches.[25]

Importantly, our identification strategy does not hinge on strong assumptions about the exogeneity of tariffs or other aggregate variables. While the estimated parameters may be shaped by general equilibrium forces, identification relies solely on cross-sectional variation across buyers within supplier–product–year cells. As a result, we do not need to specify or model the broader general equilibrium environment.

Estimation We estimate equation (4.1) via generalized method of moments (GMM),

$$\min_{\{\phi,\theta\}} \ \mathbf{g}(\phi,\theta)\ \mathbf{Z}'\ \mathbf{W}\ \mathbf{Z}\ \mathbf{g}(\phi,\theta)', \tag{4.2}$$

where $\mathbf{g}(\phi, \theta)$ stacks all moment conditions in equation (4.1) across all i–j–ℓ pairs and years and $\mathbf{W}$ is the optimal weighting matrix.[26]

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

[25] Formally, identification relies on the nonlinearity of the markup equation (2.6) in s_{ijht} and x_{ijht}. Consider moment conditions from two periods t and $t-1$: the associated derivatives with respect to (ϕ, θ) are not collinear, satisfying the full-rank condition. Similar variation across multiple matches for the same exporter in a given year (e.g., i–j–k vs. i–j–ℓ) also secures identification.

[26] Appendix E.2 presents Monte Carlo simulations based on a data-generating process that mirrors the setup in Section 2. The results confirm that our estimators is consistent.

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:

$$\phi_{ijt} = \frac{\exp\left(\mathbf{X}_{ijht}\,\boldsymbol{\kappa}\right)}{1 + \exp\left(\mathbf{X}_{ijht}\,\boldsymbol{\kappa}\right)} \in [0, 1], \tag{4.3}$$

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 relative outside option of the two, 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 (4.2) using data from 2001 to 2016. We exclude 2017 and 2018, as these years will be used to validate the model out-of-sample in Section 4.3, leveraging the tariff shocks that occurred during this period. To avoid convergence issues when ϕ is near 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 (4.3). 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 0.70 to 0.92.

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: $\widehat{\ln \frac{\phi}{1-\phi}}$	1.565		0.863	
	(0.055)		(0.043)	
Returns to scale ($\hat{\theta}$)	0.454	0.497	0.383	0.502
	(0.004)	(0.006)	(0.006)	(0.007)
Constant		4.118		1.454
		(0.428)		(0.180)
Longevity		-0.360		0.332
		(0.062)		(0.064)
Number of HS10 transactions		-0.264		-0.003
		(0.029)		(0.014)
Multiple HS10 dummy		-0.180		0.131
		(0.047)		(0.034)
Lagged outside option		-0.235		-0.230
		(0.031)		(0.030)
None	Yes	Yes	No	No
$\text{FE}_h + \text{FE}_t + \text{FE}_j$	No	No	Yes	Yes
Observations	3,120,000			
Panel C: Implied Bargaining Powers ($\hat{\phi}$)				
Mean	0.827	0.922	0.703	0.860
	(0.008)	(0.074)	(0.009)	(0.099)
Median	–	0.945	–	0.886
	–	(0.074)	–	(0.099)

Notes: This table presents model estimates based on our main estimation sample, which focuses on U.S. imports of intermediate inputs and capital goods for the period 2001-2016. Panel A reports calibrated parameters: the elasticity of demand (ν), the elasticity of costs with respect to foreign input prices (γ), and the elasticity of substitution across foreign varieties (ρ). We set $\varrho = 1$, so that $\eta = 2.5$. Panel B presents GMM estimates. Columns (1) and (3) impose a constant ϕ across bilateral pairs, while Columns (2) and (4) estimate the full vector $\boldsymbol{\kappa}$ to allow for heterogeneity in bargaining power. Specifications differ in the inclusion of fixed effects. Controls include: (i) the log of relationship longevity between exporter i and importer j; (ii) the log of the number of transactions between i and j in a given year; (iii) the log of the relative outside option, defined as the ratio of exporter i's sales to other U.S. buyers (excluding j) over importer j's purchases from other suppliers (excluding i), both in year $t-1$; and (iv) a dummy variable equal to one if the $i - j$ pair transacts in more than one HS10 product. Panel C reports the mean and median of the implied bargaining power. Standard errors are robust; those in Panel C are computed using the delta method. The set of instruments includes the number of exporters and importers at the HS10 level, as well as lagged bilateral shares (excluding the focal pair). 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).

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.[27]

The returns to scale parameter $\hat{\theta}$ is consistently estimated below one, ranging from 0.40 to 0.50 across specifications, with a preferred estimate of 0.45.[28] This implies a residual export supply elasticity between 0.25 and 0.375 for the average importer, indicating relatively steep foreign supply curves.[29] 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 $\hat{\boldsymbol{\kappa}}$, we find that the coefficients on relationship longevity, frequency of transactions, and the multiple-product indicator are highly significant, although their sign varies depending on the set of fixed effects included. By contrast, the coefficient on the relative outside option is stable across 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 our structural estimates to alternative sample definitions and model calibrations. Appendix D.3 first considers a broader sample that includes all products in the BEC classification, notably extending the baseline by adding consumption goods. We also examine an alternative sample that uses related-party trade indicators from the LFTTD instead of ORBIS. On the calibration side, we consider a lower elasticity of substitution across foreign varieties ($\rho = 5$ instead of 10) and introduce decreasing returns to scale in downstream production by setting $\varrho = 0.5$ rather than 1. This choice aligns with the estimated returns to scale on the exporter side and allows us to test the robustness of the estimates to alternative values of the importer's downstream demand elasticity, η. Across all variations, the structural estimates remain highly stable.

Implied Markups Using equation (2.6), we compute markups for all buyer–supplier–product matches given the estimated parameters and the observed distribution of market shares.

[27] These findings are consistent with evidence from related settings. Morlacco (2019) documents significant markdowns by French importers in input trade, while Atkin et al. (2024) show that Argentine importers often exercise considerable bargaining power.

[28] While standard estimates of returns to scale often cluster near one, θ here captures a short-run returns to scale elasticity, i.e., the slope of the exporter's marginal cost curve during bargaining.

[29] The implied residual supply elasticity is computed from the marginal cost slope $c'_{q_{ijh}} = \frac{1-\theta}{\theta} \cdot x_{ijh}$, using the average buyer share of 0.25 from Table 1.

The resulting markup distribution is clustered near the competitive benchmark. Our preferred estimates yield a mean markup of 0.94, with the median even closer to competitive levels. These low markups reflect the strong countervailing power of importers. The model implies that observed markups are a convex combination of oligopoly and oligopsony markups, which average 1.34 and 0.87, respectively. The large estimated bargaining power of buyers ($\hat{\phi} \approx 0.8$) shifts weight toward the oligopsony case, allowing importers to extract a substantial share of the surplus from exporters.[30]

4.3 Model Validation

We assess the model's empirical validity by testing its ability to predict both the level and heterogeneity of price and quantity changes following the 2017–2018 tariff increases. These moments were not targeted in the estimation.

Model-predicted price changes are computed as:

$$\widehat{\Delta \ln p_{ijht}} = \Phi_{ijht}(s_{ijht}, x_{ijht} \mid \hat{\boldsymbol{\Theta}}) \cdot \Delta \ln(1 + \tau_{cht}), \tag{4.4}$$

where τ_{cht} denotes the ad-valorem tariff on product h from country c, and Φ_{ijht} is the model-implied pass-through elasticity defined in equation (2.11), which depends on bilateral shares and the estimated parameter vector $\hat{\boldsymbol{\Theta}}$.

Although the model is primarily designed to explain prices, it embeds a demand-driven allocation rule that links prices to quantities (and sales) via the importer's demand curve. This mapping imposes a specific quantity-setting structure, enabling a direct comparison between predicted price responses and observed adjustments in quantities and trade values. Predicted quantity changes are given by $\widehat{\Delta \ln q_{ijht}} = -\hat{\varepsilon}_{ijht} \cdot \widehat{\Delta \ln p_{ijht}}$, where $\hat{\varepsilon}_{ijht}$ denotes the match-specific residual demand elasticity implied by the model. For sales, the corresponding mapping is $\widehat{\Delta \ln r_{ijht}} = -(1 - \hat{\varepsilon}_{ijht}) \cdot \widehat{\Delta \ln p_{ijht}}$.

4.3.1 Price Predictions

We begin with price outcomes. Table 5 compares tariff pass-through elasticities in the data (Panel A) and in the model (Panel B), in terms of average effects and their heterogeneity with respect to buyer and supplier shares. Columns (1) and (2) adopt a baseline specification with product–time and country–sector fixed effects, while Columns (3) and (4) adopt the more demanding specification with buyer–time and country–time fixed effects.

[30] The effective bargaining weight ω_{ijh} averages 0.77 (standard deviation 0.05), slightly below $\hat{\phi}$. This indicates that network effects, on average, dampen the importers' effective bargaining power relative to ϕ, although the gap is small.

Table 5: Price Responses and Relationship Heterogeneity: Data vs. Model

Panel A: Data

	(1)	(2)	(3)	(4)
$\Delta \ln(1+\tau_{cht})$	-0.151	-0.066	-0.223	-0.163
	(0.093)	(0.093)	(0.109)	(0.107)
$\Delta \ln(1+\tau_{cht}) \cdot s_{ijht-1}$		0.054		0.029
		(0.072)		(0.157)
$\Delta \ln(1+\tau_{cht}) \cdot x_{ijht-1}$		-0.403		-0.271
		(0.113)		(0.135)
R-squared	0.04	0.04	0.31	0.31

Panel B: Model

	(1)	(2)	(3)	(4)
$\Delta \ln(1+\tau_{cht})$	-0.248	-0.144	-0.249	-0.136
	(0.008)	(0.009)	(0.010)	(0.011)
$\Delta \ln(1+\tau_{cht}) \cdot s_{ijht-1}$		0.104		0.091
		(0.011)		(0.010)
$\Delta \ln(1+\tau_{cht}) \cdot x_{ijht-1}$		-0.475		-0.486
		(0.026)		(0.029)
R-squared	0.32	0.46	0.50	0.59
$\text{FE}_{ht} + \text{FE}_{cs}$	Yes	Yes	No	No
$\text{FE}_{ht} + \text{FE}_{ct} + \text{FE}_{jt}$	No	No	Yes	Yes
Observations	249,000			

Notes: This table reports the pass-through of tariffs to duty-exclusive prices at the exporter–importer–product level. Panel A presents reduced-form estimates from the data. Panel B shows corresponding pass-through estimates generated by the model. Columns (2)–(4) interact tariff changes with lagged supplier share (s_{ijht-1}) and lagged buyer share (x_{ijht-1}). Columns (1) and (2) use baseline fixed effects ($\text{FE}_{ht} + \text{FE}_{cs}$), while Columns (3) and (4) employ a more stringent specification with product–year, country–year, and buyer–year fixed effects ($\text{FE}_{ht} + \text{FE}_{ct} + \text{FE}_{jt}$). 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).

The model predicts an average pass-through elasticity on duty-exclusive prices of -0.25, which translates to a 75% pass-through rate (Columns (1) and (3) of Panel B). These values fall within the 95% confidence intervals of the corresponding reduced-form estimates in the data (Columns (1) and (3) of Panel A).

Columns (2) and (4) of Panel A, consistent with Column (6) of Table 3, show that pass-through declines significantly with buyer share (x_{ijht}) but exhibits no robust relationship with supplier share (s_{ijht}). The model closely replicates these patterns: as shown in Panel B, pass-through declines steeply in x_{ijht} and increases mildly in s_{ijht}, with magnitudes comparable to those in the data. This alignment supports our interpretation of Table 3 as consistent with a low-θ, high-ϕ environment, which is also consistent with our structural estimates.

4.3.2 Goodness-of-fit Test

Having shown that the model replicates average and heterogeneous price responses to tariff shocks, we formally test its predictive performance by evaluating how well model-implied price changes explain observed variation. This exercise complements Table 5 by providing a direct measure of goodness of fit relative to standard alternatives.

Specifically, we estimate:

$$\Delta \ln p_{ijht} = \beta \, \widehat{\Delta \ln p_{ijht}} + \mathbf{FE} + u_{ijht}, \tag{4.5}$$

where $\Delta \ln p_{ijht}$ is the observed change in the duty-inclusive price for product h between exporter i and importer j, and $\widehat{\Delta \ln p_{ijht}}$ is the corresponding model-predicted change.

To benchmark the results, we compare the full model to three nested alternatives that sequentially shut down bilateral bargaining ($\phi = 0$), decreasing returns to scale ($\theta = 1$), or both. The fully restricted case ($\phi = 0$, $\theta = 1$) corresponds to a standard Nash-Bertrand model with constant marginal costs (e.g., Dhyne et al., 2022); the intermediate case with $\phi > 0$ and $\theta = 1$ mirrors the bargaining framework of Gopinath and Itskhoki (2011), which overlooks the heterogeneous buyer dimension. In all cases, we use the same parameter values from Column (1) of Table 4, without re-estimating.[31]

A key challenge is that observed price changes may reflect shocks unrelated to tariffs, increasing the risk of rejecting a valid model for reasons unrelated to its tariff predictions. In addition, while the model uses statutory tariff changes, observed duty-inclusive prices re-

[31] As discussed in Section 4.1, θ is not identified when $\phi = 0$. Moreover, fixing $\theta = 1$ and estimating equation (4.5) using the value of ϕ jointly obtained with θ from equation (4.2) yields an upper bound on the value of $\hat{\beta}$ attainable with a re-estimated ϕ. Appendix E.3 provides formal proof and supporting simulation evidence.

FIGURE 3: IV-Based Goodness-of-Fit Test

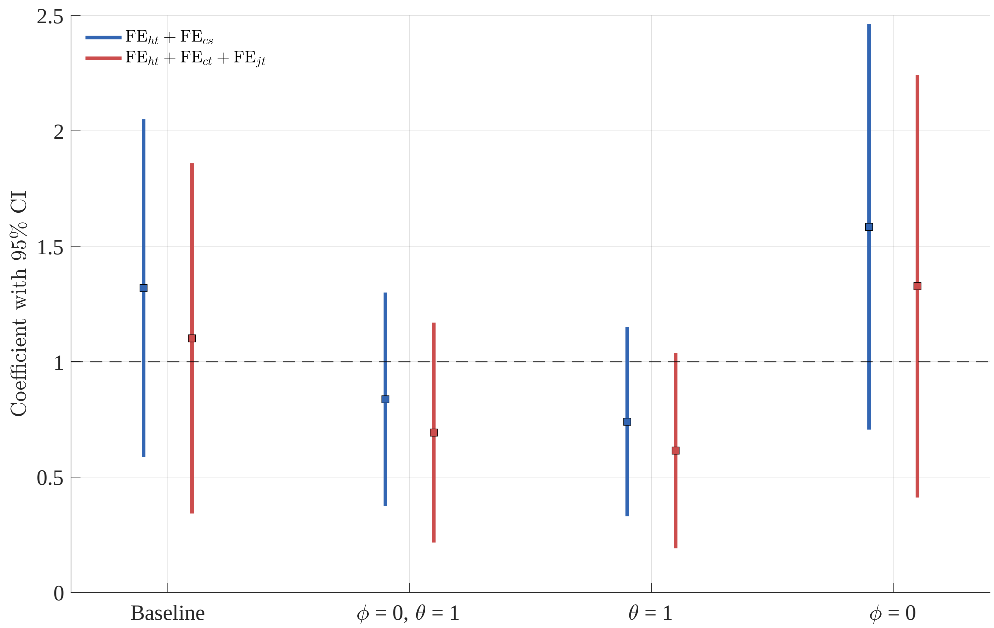

Notes: Each point reports the coefficient from an IV regression of observed log price changes on model-predicted changes $\widehat{\Delta \ln p_{ijht}}$, using statutory tariffs as instruments. Lines show 95% confidence intervals. Blue and red denote regressions with $\mathrm{FE}_{ht} + \mathrm{FE}_{cs}$ and $\mathrm{FE}_{ht} + \mathrm{FE}_{ct} + \mathrm{FE}_{jt}$, 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).

flect actual duties paid, potentially introducing measurement error. To address both issues, we follow the logic of the IV-based goodness of fit test in Adão et al. (2023) and estimate equation (4.5) using two-stage least squares, instrumenting $\widehat{\Delta \ln p_{ijht}}$ with statutory tariff changes. This isolates variation in predicted prices that is directly attributable to tariff shocks, which is the one targeted by the model. Under the null that the model accurately captures pass-through, the IV coefficient $\hat{\beta}$ should equal one.

Figure 3 presents the results across the four model variants, each estimated under two alternative fixed effects specifications. Blue points and lines correspond to regressions with product–time and country–sector fixed effects, while red ones use the more stringent specification with product–time, country–time, and buyer–time fixed effects. In all cases, the baseline model provides the best fit: the estimated coefficient is very close to one and not statistically different from it under both specifications.

Models with constant returns to scale ($\theta = 1$), whether or not they include bargaining, perform noticeably worse, although their coefficients are not rejected at conventional significance levels. These variants lack a cost channel, and while the case with $\phi = 0$ features strategic complementarities that generate some degree of incomplete pass-through, this mechanism alone does not produce sufficient variation to match the data. Similarly, the fourth model with decreasing returns but no bargaining ($\phi = 0$, $\theta < 1$), which was previously rejected based on price-level evidence (Table 2), also underperforms relative to the baseline model. In this case, markup responses driven by strategic complementarities amplify cost-based adjustments, but in a way that also fails to replicate the observed pass-through patterns.[32]

Overall, these results suggest that cost adjustments, along with weak strategic complementarities, are necessary to account for the empirical evidence. This is consistent with a setting in which oligopsony forces are dominant.

4.3.3 Quantity Predictions

Next, we assess the model's ability to predict changes in bilateral quantities. Table D.7 in Appendix D.4 reports relationship-level quantity responses to tariff changes. Columns (3)–(4) of Panel A show that tariff increases reduce traded volumes, with larger declines under more stringent fixed effects. However, the interaction terms with supplier and buyer shares are statistically imprecise, suggesting inconclusive evidence of heterogeneity in the data.[33]

Panel B shows that the model generates sizable average quantity declines and predicts heterogeneity across relationships. In particular, quantity responses become less negative with higher buyer shares, consistent with the model's allocative logic: stronger buyers face smaller price increases and thus smaller quantity reductions. In contrast, the positive interaction with supplier share is not explained by price adjustments because the model predicts little variation in pass-through along this dimension. Instead, it reflects the curvature of the importer's demand curve embedded in the model.

Table D.8 in Appendix D.4 evaluates model fit for quantities (Panel A) and sales (Panel B) across the four alternative parameterizations discussed above. Since all models share the same demand system, performance differences reflect variation in price predictions rather than differences in quantity mechanisms. While all specifications are formally rejected, the baseline model and the specification with decreasing returns to scale perform best, mirroring their superior performance in prices.

Despite differences in magnitude and precision, the model captures the broad directional

[32] Appendix D.4 (Table D.6) presents additional robustness checks using alternative calibrations, including $\rho = 5$ instead of 10 and $\varrho = 0.5$ instead of 1. The model's predictive performance remains stable across these variations, lending further support to our main findings.

[33] In robustness exercises (not shown), the signs on the interaction coefficients with supplier and buyer shares occasionally flip, though they largely remain statistically insignificant. We therefore conclude that the data do not offer robust evidence of heterogeneous quantity responses by relationship structure.

patterns in quantity responses, lending support to its core allocative mechanism. However, its weaker quantitative fit suggests that additional forces beyond price-based allocation shape the observed quantity adjustments. The model's strong performance on prices highlights this asymmetry and points to potential gains from extending the framework to incorporate supply-side responses.

5 Aggregate Implications of Bargaining in Firm-to-Firm Trade

This section examines the impact of the Trump tariffs on aggregate import prices through the lens of our bargaining model. We begin by situating our analysis within the broader literature on tariff pass-through. We then use the model to simulate the aggregate effects of the tariffs and assess the mechanisms driving price adjustments.

5.1 Comparison with Existing Pass-Through Estimates

Several recent studies find near-complete pass-through of the 2018 Trump tariffs to U.S. import prices (e.g., Fajgelbaum et al., 2020; Amiti et al., 2019b, 2020). These analyses primarily rely on product-level data. In contrast, our match-level estimates indicate incomplete pass-through. For instance, in a standard specification similar to those used in the literature, Column (1) of Table 5 (Panel A) reports a pass-through elasticity on duty-exclusive prices of -0.15, implying an 85% pass-through rate.

This divergence partly reflects differences in sample composition. As discussed in Section 3.3, our estimates are based on repeated firm-to-firm relationships where bilateral prices can be observed over time. They therefore capture within-relationship price changes net of compositional shifts across firms or products, rather than on changes in product-level unit values. In addition, we restrict attention to arm's-length transactions involving intermediate goods and suppliers with two or more buyers. Within this sample, we estimate pass-through rates ranging from 78 to 85%, depending on the fixed effects used, as shown in Table 5. The model's predicted pass-through closely matches these empirical patterns.

To further assess the role of sample composition, Table D.9 in Appendix D.5 reports pass-through estimates across alternative sample definitions. Expanding the baseline to include single-buyer relationships raises the estimate to approximately 83%. Adding matches involving related parties, energy goods, or extreme price changes increases it to roughly 90%. The most inclusive specification, which further incorporates final consumption goods, yields pass-through elasticities between 93% and 95%, depending on the fixed effects. Although these broader samples imply higher pass-through, they still fall short of full pass-through, in contrast to product-level studies.[34]

These product-level estimates are based on aggregated unit values that combine within-match price changes with shifts in the composition of transactions, particularly the inclusion of one-off (spot) exchanges. Such transactions may differ systematically from repeated firm-to-firm relationships. In particular, incomplete pass-through may be a feature of relational trade, where prices reflect bilateral bargaining and upward-sloping residual supply. Spot exchanges, by contrast, are more likely to reflect marginal cost pricing and thus exhibit full pass-through. While speculative, this interpretation is consistent with Cajal-Grossi et al. (2023), who find lower markups and more competitive pricing in spot relationships.

5.2 Tariffs and Aggregate Import Prices

We use the estimated model to assess the impact of the 2018 tariffs on aggregate import prices and decompose the underlying contributions of markup and marginal cost changes. A full evaluation of the tariffs' macroeconomic effects would require a general equilibrium framework that incorporates changes in expenditure, wages, export prices, and retaliation, as in Fajgelbaum et al. (2020). While such elements are essential for welfare analysis, they lie beyond the scope of our model.

Instead, we focus on one central component of tariff incidence: the change in aggregate import prices within ongoing firm-to-firm relationships. We compute model-implied bilateral price changes using equation (4.4) and estimate aggregate pass-through by regressing these predicted changes on the tariff shock, weighting each observation by its initial-period import value. The approach connects the model's micro-level predictions to aggregate outcomes and provides a framework for decomposing the contribution of distinct adjustment channels.

To that end, we separately isolate the roles of markup and marginal cost adjustments. For each mechanism, we compute the model-implied price changes while holding the other channel constant, and re-estimate the weighted regression. The predicted price change due solely to markup adjustment is defined as:

$$\widehat{\Delta^{\Gamma} \ln p_{ijht}} = \Phi^{\Gamma}_{ijht}(s_{ijht},\, x_{ijht} \mid \hat{\boldsymbol{\Theta}}) \cdot \Delta \ln T_{cht}, \quad \text{where} \quad \Phi^{\Gamma}_{ijht} \equiv \frac{1}{1+\Gamma_{ij}}.$$

[34] Using the monthly LFTTD data, we show in Table D.10 in Appendix D.5 that pass-through remains incomplete in product-level regressions when restricting to the firms and products used in our baseline sample, confirming that this is a feature of repeated firm-to-firm relationships.

Table 6: Aggregate Tariff Pass-Through and Decomposition

	Baseline FE	Stringent FE
	(1)	(2)
Panel A: Aggregate Passthrough (%)		
Agg. pass-through elasticity	70.7	64.6
Cost channel only: $1/(1+\Lambda_{ij})$	69.3	62.7
Markup channel only: $1/(1+\Gamma_{ij})$	97.1	88.0
Panel B: Variance Decomposition of $\Lambda_{ij}+\Gamma_{ij}$		
Cost Elasticity: Λ_{ij}	1.01	0.99
Markup Elasticity: Γ_{ij}	-0.01	0.01

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, while the "Markup channel only" row sets cost elasticities 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).

Similarly, the predicted price change driven solely by cost adjustment is given by:

$$\widehat{\Delta^{\Lambda}\ln p_{ijht}} = \Phi^{\Lambda}_{ijht}(s_{ijht},\, x_{ijht} \mid \hat{\boldsymbol{\Theta}}) \cdot \Delta\ln T_{cht}, \quad \text{where} \quad \Phi^{\Lambda}_{ijht} \equiv \frac{1}{1+\Lambda_{ij}}.$$

Panel A of Table 6 reports the aggregate tariff pass-through predicted by the model, which incorporates both markup and marginal cost adjustments: 71% under the baseline specification (Column (1)) and 65% under the alternative using product–time, country–time, and buyer–time fixed effects (Column (2)). These estimates are lower than the average pass-through of 75% obtained from unweighted regressions (Table 5), reflecting the more incomplete pass-through associated with larger, high-share buyers who account for a greater share of aggregate trade.

Nearly all of the predicted price response is driven by the cost channel. When markup elasticities are shut down (i.e., setting $\Gamma_{ij}=0$), the model yields pass-through rates between 69 and 63%, closely matching the full-model estimates. By contrast, shutting down the cost elasticities results in much higher pass-through, ranging from 88 to 97%.

Panel B offers an alternative approach to quantify the contribution of each channel by decomposing the variance of the term $\Lambda_{ij} + \Gamma_{ij}$, which governs the pass-through elasticity $1/(1+\Lambda_{ij}+\Gamma_{ij})$. This decomposition attributes the cross-sectional variation in pass-through to cost and markup elasticities, with their shares summing to one by construction. The results reinforce the earlier findings: nearly all of the variation is explained by the cost elasticity Λ_{ij}, underscoring its dominant role in shaping price responses.

Together, the results suggest a clear conclusion: tariff pass-through is incomplete, reflecting the combined effects of cost-side adjustments and bargaining. The underlying mechanism is intuitive. When tariffs reduce demand from large U.S. buyers, exporters move down their marginal cost curves, dampening the price impact of the shock. Although markups adjust endogenously, the opposing forces of strategic complementarities and substitutabilities limit their quantitative contribution to aggregate price change. This is a robust implication of our framework. As illustrated in Figure 2, when returns to scale are decreasing ($\theta < 1$) and buyer power is high (large ϕ), most of the variation in pass-through across the $(s_{ij},\, x_{ij})$ space is driven by cost adjustments, with a comparatively smaller role for markup responses.

5.3 The Importance of Bargaining in Firm-to-Firm Trade

Although bargaining plays a central role in our model, it appears to have limited influence on tariff pass-through, which primarily reflects cost-side adjustments. This is not because bargaining is unimportant, but because it endogenously mutes the markup channel. When buyer power is strong (ϕ high), the markup elasticity (Γ_{ij}) tends toward zero. As a result, pass-through is driven almost entirely by cost adjustments and appears relatively unresponsive to markup variation. In this sense, the weak role of the markup channel in shaping pass-through is itself a consequence of strong importer market power.

In addition, bargaining remains essential to understand price levels, markups, and the welfare consequences of market power. As shown in Table 2, two-sided market power is critical to explaining the cross-sectional variation in prices and markups. Moreover, as discussed in Section 2.5, ϕ governs how markup dispersion maps into misallocation. When exporters hold bargaining power, markup heterogeneity leads to underproduction. But when bargaining power shifts toward importers, the same dispersion can improve efficiency by reallocating output toward lower-cost suppliers. Ignoring bargaining would therefore yield misleading conclusions about the allocative effects of firm-to-firm trade.

6 Conclusions

Firm-to-firm relationships are a central feature of international trade. These relationships often involve market power on both sides, with prices determined through bilateral negotiations rather than market-clearing conditions. We develop a framework that departs from standard models of price-taking buyers by allowing both importers and exporters to influence price formation through bargaining. This approach yields analytical and empirical tools to study how market power and network frictions shape prices, markups, and the transmission of shocks in firm-to-firm trade.

Using transaction-level data, we show that U.S. importers wield substantial bargaining power and face upward-sloping residual supply, consistent with oligopsonistic behavior. Within ongoing firm-to-firm relationships, tariff pass-through during the 2018 trade war was incomplete, with exporters absorbing much of the tariff through cost-side adjustments. This muted price response is primarily driven by cost-side adjustments: while markups do adjust, their contribution to pass-through is limited because strong buyer power endogenously flattens the markup elasticity. In this sense, the limited role of markups in shaping tariff responses is itself a consequence of strong importer power.

Although not the primary focus of our analysis, the results highlight the importance of bargaining for understanding the allocative implications of market power in firm-to-firm trade. In our model, when supplier power dominates, markup dispersion reflects inefficiencies and underproduction. By contrast, when buyer power is strong, the same dispersion can lead to efficient reallocation and higher aggregate output. These findings suggest that markup heterogeneity does not uniformly signal misallocation, and that its welfare consequences depend critically on the distribution of bargaining power. Extending the framework to study these welfare implications more formally remains a promising avenue for future work.

More broadly, our study offers a foundation for analyzing how market power shapes the incidence and transmission of shocks through the price channel. While we focus on detailed price-setting within firm-to-firm relationships, Acemoglu and Tahbaz-Salehi (2025) emphasize the general equilibrium implications of market power in production networks in shaping aggregate fluctuations. A promising direction for future research is to integrate these perspectives by combining micro-level bargaining dynamics with macro-level spillovers to study the broader implications of buyer and supplier power, including in domestic supply chains.

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 US tariffs? A longer-term perspective," in *AEA Papers and Proceedings*, vol. 110, 541–46.

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.

——— (2020): "Conceptual Aspects of Global Value Chains," *The World Bank Economic Review*.

Antràs, P. and D. Chor (2013): "Organizing the global value chain," *Econometrica*, 81, 2127–2204.

Antràs, P. and R. W. Staiger (2012): "Offshoring and the role of trade agreements," *American Economic Review*, 102, 3140–83.

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, Markdowns, and Bargaining in a Vertical Supply Chain," *Available at SSRN*.

Barrot, J. N. and J. Sauvagnat (2016): "Input Specificity and the Propagation of Idiosyncratic Shocks in Production," *Quarterly Journal of Economics*, 131, 1543 – 1592.

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.

Blaum, J., C. Lelarge, and M. Peters (2018): "The gains from input trade with heterogeneous importers," *American Economic Journal: Macroeconomics*, 10, 77–127.

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. Limao, 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.

Burstein, A. and G. Gopinath (2015): "International Prices and Exchange Rates," in *Handbook of International Economics*, vol. 4, 391–451.

Cajal-Grossi, J., R. Macchiavello, and G. Noguera (2023): "Buyers' sourcing strategies and suppliers' markups in Bangladeshi garments," *The Quarterly Journal of Economics*, 138, 2391–2450.

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.*

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

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.*

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.*

Grossman, G. M., E. Helpman, and S. J. Redding (2024): "When tariffs disrupt global supply chains," *American Economic Review*, 114, 988–1029.

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, and F. Pisch (2021): "Firm heterogeneity and imperfect competition in global production networks," *Available at SSRN.*

Kamal, F. and R. Monarch (2018): "Identifying Foreign Suppliers in U.S. Import Data," *Review of International Economics*, 26, 117–139.

Macchiavello, R. and A. Morjaria (2015): "The value of relationships: evidence from a supply shock to Kenyan rose exports," *American Economic Review*, 105, 2911–45.

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

Monarch, R. (2022): "''It's Not You, It's Me": Prices, Quality, and Switching in US-China Trade Relationships," *Review of Economics and Statistics*, 104, 909–928.

Morlacco, M. (2019): "Market Power in Input Markets : Theory and Evidence from French Manufacturing," *mimeo*.

Noll, R. G. (2005): "Buyer power and economic policy," *Antitrust LJ*, 72, 589.

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

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.

Stole, L. A. and J. Zwiebel (1996): "Intra-firm bargaining under non-binding contracts," *The Review of Economic Studies*, 63, 375–410.

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

This appendix derives the key equations from the main text: the bilateral markup μ_{ij} (equation (2.6)) and the bilateral pass-through elasticity Φ_{ij} (equation (2.11)).

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 (2.4):

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

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

$$q_j^f = \gamma c_j q_j \left(p_j^f \right)^{-1}, \tag{A.1}$$

$$q_j^d = (\varrho - \gamma) c_j q_j \left(p_j^d \right)^{-1}, \tag{A.2}$$

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

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

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 $\frac{d \ln c_j}{d \ln p_j^f}$, namely, the elasticity of the marginal cost c_j with respect to p_j^f. To find this elasticity, we first use the demand function downstream to write q_j as

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

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

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

which implies

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

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

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

and equation (A.1), we find equation (2.9):

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

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:

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

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:

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

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

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

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

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

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

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

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:

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

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

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

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

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

where c_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:

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

The outside profit of firm j under a failed negotiation is

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

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

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

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

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

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:

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

where the weighting factor is

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

and where

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

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} co-moves with the exporter's and importer's bilateral market shares, s_{ij} and x_{ij}.

Part (i): μ_{ij} and Exporter's supplier share s_{ij} To study how μ_{ij} responds to the exporter's supplier share s_{ij}, note that the oligopoly markup $\mu_{ij}^{\text{oligopoly}}$ is strictly increasing in s_{ij} under standard assumptions ($\eta < \rho < \infty$). Differentiating the overall markup μ_{ij} in equation (A.12) with respect to s_{ij} yields:

$$\frac{\partial \mu_{ij}}{\partial s_{ij}} = (1-\omega_{ij}) \cdot \underbrace{\frac{\partial \mu_{ij}^{\text{oligopoly}}}{\partial s_{ij}}}_{>0} + \frac{\partial \omega_{ij}}{\partial s_{ij}} \cdot \underbrace{\left(\mu_{ij}^{\text{oligopsony}} - \mu_{ij}^{\text{oligopoly}} \right)}_{<0}.$$

The first term is strictly positive when $\omega_{ij} < 1$ (i.e., $\phi < 1$), reflecting the direct effect of s_{ij} on the oligopoly markup. The second term captures how changes in s_{ij} affect the bargaining weight ω_{ij} and thus the relative influence of the oligopsony markdown. Although ω_{ij} is hump-shaped in s_{ij}, increasing at low values and decreasing at high values, its sensitivity to s_{ij} is limited in most of the parameter space. As a result, $\frac{\partial \omega_{ij}}{\partial s_{ij}}$ is typically small, and the first (positive) term generally dominates.

In contrast, when $\phi = 1$ (full importer bargaining power), the bilateral markup reduces to the oligopsony markdown: $\mu_{ij} = \mu_{ij}^{\text{oligopsony}}$, and $\omega_{ij} = 1$. In this case,

$$\frac{\partial \mu_{ij}}{\partial s_{ij}} = 0.$$

Thus, we have:

$$\frac{\partial \mu_{ij}}{\partial s_{ij}} = \begin{cases} \gtreqless 0 & \text{if } \phi < 1 \text{ (typically positive)}, \\ = 0 & \text{if } \phi = 1. \end{cases}$$

Thus, μ_{ij} can increase with s_{ij} in our theory only when $\phi < 1$, which proves the first part of Proposition 1.

Part (ii): μ_{ij} and Importer's buyer share x_{ij} The effect of x_{ij} on μ_{ij} is limited to the oligopsony markdown term and is therefore easier to characterize. Differentiating the bilateral markup with respect to x_{ij} yields:

$$\frac{\partial \mu_{ij}}{\partial x_{ij}} = \omega_{ij} \cdot \underbrace{\frac{\partial \mu_{ij}^{\text{oligopsony}}}{\partial x_{ij}}}_{<0 \text{ if } \theta<1}.$$

The markdown $\mu_{ij}^{\text{oligopsony}}$ decreases with x_{ij} when marginal cost is increasing (i.e., $\theta < 1$), as larger buyers elicit stronger cost reductions and thus negotiate lower prices. This implies:

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

Therefore, observing a negative slope of μ_{ij} with respect to the importer's share x_{ij} provides evidence of both importer bargaining power ($\phi > 0$ so that $\omega_{ij} > 0$) and decreasing returns to scale ($\theta < 1$). This completes the proof of Proposition 2. □

A.3 Proof of Proposition 3

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

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

where c_i and μ_{ij} are as in equations (2.1) and (2.6), 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 $\Lambda_{ij} \equiv -\frac{d\ln c_i}{d\ln p_{ij}}$ are the partial *markup* and *cost* elasticities, respectively.

The Cost Elasticity

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

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

It immediately follows that:

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

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

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

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

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

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 (2.6) and differentiating, we obtain:

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

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:

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

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

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

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

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

which implies:

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

with

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

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

Omega Elasticity

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

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

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:

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

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:

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

$\square$

A.4 Proof of Proposition 4

We now prove Proposition 4, which characterizes the equilibrium comovement between pass-through Φ_{ij} and the bilateral market shares x_{ij} and s_{ij}.

Comparative Statics with Respect to x_{ij}— We begin by analyzing how pass-through Φ_{ij} responds to the importer's buyer share x_{ij}. As discussed in the main text, Φ_{ij} depends on x_{ij} through two channels: (i) the oligopsony markdown elasticity $\Gamma_{ij}^{\text{oligopsony}}$, which is non-monotonic (U-shaped) in x_{ij}, and (ii) the cost elasticity Λ_{ij}, which increases linearly in x_{ij} when $\theta < 1$.

When $\theta = 1$, both elasticities are zero: $\Gamma_{ij}^{\text{oligopsony}} = \Lambda_{ij} = 0$. Therefore, pass-through is constant in x_{ij}:

$$\left.\frac{\partial \Phi_{ij}}{\partial x_{ij}}\right|_{\theta=1} = 0.$$

When $\theta < 1$, however, both elasticities are active. The pass-through elasticity can be written as:

$$\Phi_{ij} = \frac{1}{1 + \Gamma_{ij} + \Lambda_{ij}} = \frac{1}{1 + \cdots + \omega_{ij} \cdot \frac{\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} \cdot \Gamma_{ij}^{\text{oligopsony}} + \Lambda_{ij}},$$

where "$\cdots$" denotes terms that do not depend on x_{ij}. Differentiating with respect to x_{ij}:

$$\frac{\partial \Phi_{ij}}{\partial x_{ij}} = -\frac{1}{(1 + \Gamma_{ij} + \Lambda_{ij})^2} \cdot \frac{\partial}{\partial x_{ij}} \left(\omega_{ij} \cdot \frac{\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} \cdot \Gamma_{ij}^{\text{oligopsony}} + \Lambda_{ij} \right).$$

Let:

$$T_1 \equiv \omega_{ij} \cdot \frac{\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} \cdot \Gamma_{ij}^{\text{oligopsony}}, \quad T_2 \equiv \Lambda_{ij}.$$

We can expand these as:

$$\begin{aligned} T_1 &= \omega_{ij} \cdot \left(\frac{(1 - x_{ij})^{\frac{1}{\theta} - 1} - \mu_{ij}^{\text{oligopsony}}}{\mu_{ij}} \right) (1 - x_{ij}) \varepsilon_{ij}, \\ T_2 &= \frac{1 - \theta}{\theta} \cdot x_{ij} \cdot \varepsilon_{ij}. \end{aligned}$$

While T_1 is non-monotonic in x_{ij}, T_2 increases linearly in x_{ij} and dominates T_1 for all x_{ij}. This can be verified numerically; for instance, when $\theta = \frac{1}{2}$:

$$T_1 + T_2 = x_{ij} \cdot \varepsilon_{ij} \left(1 - \frac{\omega_{ij}}{\mu_{ij}} \cdot \left(\frac{5}{2} - \frac{3}{2} x_{ij} \right) \right),$$

which increases in x_{ij} as the term in parentheses remains positive.

Hence, the term in the numerator of $\frac{\partial \Phi_{ij}}{\partial x_{ij}}$ increases in x_{ij}, implying:

$$\operatorname{sign}\left(\frac{\partial \Phi_{ij}}{\partial x_{ij}} \right) = -\operatorname{sign}\left(\frac{\partial \Lambda_{ij}}{\partial x_{ij}} \right),$$

so that:

$$\frac{\partial \Phi_{ij}}{\partial x_{ij}} \leq 0.$$

This proves that, under $\theta < 1$, pass-through decreases in x_{ij}, a robust and testable implication of the model.

□

Comparative Statics with Respect to s_{ij}— We now turn to how pass-through Φ_{ij} varies with the exporter's supplier share s_{ij}. Unlike the case of x_{ij}, this relationship is more complex, as all components, namely, $\Gamma_{ij}^{\text{oligopoly}}$, $\Gamma_{ij}^{\text{oligopsony}}$, Γ_{ij}^{ω}, and Λ_{ij}, may vary with s_{ij}. Specifically,

$$\Phi_{ij} = \frac{1}{1+\Gamma_{ij}+\Lambda_{ij}} = \frac{1}{1+\frac{(1-\omega_{ij})\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopoly}} + \frac{\omega_{ij}\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopsony}} + \left(1-\frac{\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\right)\Gamma_{ij}^{\omega} + \Lambda_{ij}}.$$

Differentiating with respect to s_{ij} gives:

$$\begin{aligned}\frac{\partial \Phi_{ij}}{\partial s_{ij}} = -\frac{\partial}{\partial s_{ij}}\Bigg(&\frac{(1-\omega_{ij})\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopoly}} + \frac{\omega_{ij}\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopsony}} \\ &+ \left(1-\frac{\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\right)\Gamma_{ij}^{\omega} + \Lambda_{ij}\Bigg) \cdot (1+\Gamma_{ij}+\Lambda_{ij})^{-2}.\end{aligned}$$

When $\theta = 1$, both $\Gamma_{ij}^{\text{oligopsony}} = 0$ and $\Lambda_{ij} = 0$, and the expression simplifies to:

$$\frac{\partial \Phi_{ij}}{\partial s_{ij}} = -\frac{\partial}{\partial s_{ij}}\left(\frac{(1-\omega_{ij})\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopoly}} + \left(1-\frac{\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\right)\Gamma_{ij}^{\omega}\right) \cdot (1+\Gamma_{ij})^{-2}.$$

To build intuition, consider the limit $s_{ij} \to 1$, where $\omega_{ij} \to \phi$ and $\Gamma_{ij}^{\omega} \to 0$. Then:

$$\frac{\partial \Phi_{ij}}{\partial s_{ij}} \propto \begin{cases} -\frac{\partial}{\partial s_{ij}}\left(\frac{(1-\phi)\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopoly}} + \frac{\phi\mu_{ij}^{\text{oligopsony}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopsony}} + \Lambda_{ij}\right) & \text{if } \theta < 1, \\ -\frac{\partial}{\partial s_{ij}}\left(\frac{(1-\phi)\mu_{ij}^{\text{oligopoly}}}{\mu_{ij}}\Gamma_{ij}^{\text{oligopoly}}\right) & \text{if } \theta = 1. \end{cases}$$

The term involving $\Gamma_{ij}^{\text{oligopoly}}$ is hump-shaped in s_{ij}, while the remaining components, particularly Λ_{ij}, decline monotonically with s_{ij} due to $\frac{\partial \varepsilon_{ij}}{\partial s_{ij}} = -(\rho - \eta) < 0$. Quantitatively, these latter terms dominate in most parameter ranges.

Thus, we have:

$$\text{sign}\left(\frac{\partial \Phi_{ij}}{\partial s_{ij}}\right) \propto \begin{cases} (+) & \text{if } \theta < 1 \text{ or } s_{ij} \text{ sufficiently large,} \\ (-) & \text{if } \theta = 1 \text{ and } s_{ij} \text{ sufficiently small.} \end{cases}$$

Due to the non-monotonicity and interaction of multiple channels, the relationship between Φ_{ij} and s_{ij} does not yield a clean, general prediction. While pass-through may increase with s_{ij} under decreasing returns to scale and sufficient buyer power, the sign of this relationship depends on the relative strength of strategic complementarities, oligopsony markdowns, and cost elasticity. As a result, we omit this relationship from Proposition 3 and focus instead on the more robust prediction involving x_{ij}.

$\square$

B Theory: Discussion and Extensions

B.1 Efficient Bargaining

In the efficient bargaining setup, the importer and exporter negotiate over a two-part tariff (p, q) by maximizing a Nash product of generalized firm-specific gains from trade (GFT):

$$\max_{p,q} \left[GFT_{ij}^{i}(p,q)\right]^{1-\phi} \left[GFT_{ij}^{j}(p,q)\right]^{\phi} \quad \text{s.t.} \quad GFT_{ij}^{i} \geq 0, \quad GFT_{ij}^{j} \geq 0$$

where:

$$\begin{aligned} GFT_{ij}^{i}(p_{ij}, q_{ij}) &\equiv \pi_i(p) - \tilde{\pi}_{i(-j)} = p_{ij}q_{ij} - \left(\theta c_i q_i - \theta \widetilde{c_i q_i}\right), \\ GFT_{ij}^{j}(p_{ij}, q_{ij}) &\equiv \pi_j(p) - \tilde{\pi}_{j(-i)} = \left(p_j q_j - \tilde{p}_j \tilde{q}_j\right) - p_{ij}q_{ij}. \end{aligned}$$

This formulation corresponds to a setting in which the firm pair first selects the input quantity q_{ij} to maximize joint surplus, and then negotiates over the price p_{ij} to determine how surplus is split.

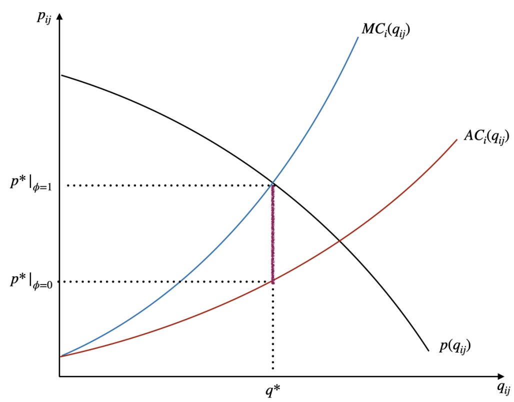

Figure B.1: Efficient Bargaining

The efficient quantity solves:

$$\begin{aligned} q_{ij}^{*}: &\quad \arg\max_{q} \left[GFT_{ij}^{i}(q) + GFT_{ij}^{j}(q)\right] \\ : &\quad \frac{d(p_j q_j)}{dq_{ij}} = \frac{d(\theta c_i q_i)}{dq_{ij}} \end{aligned}$$

That is, the efficient quantity q_{ij}^{*} equates the exporter's marginal cost to the importer's marginal revenue, similar to the outcome under vertical integration. Solving yields:

$$q_{ij}^* : \quad p_{ij} = c_i(q_{ij}^*),$$

so that the negotiated price equals the exporter's marginal cost.

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

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

where $\phi \in (0,1)$ denotes the importer's bargaining power. The terms $\Delta R_j(q_{ij})$ and $\Delta C_i(q_{ij})$ represent the additional revenue for the importer and the additional cost for the exporter attributable to the match:

$$\begin{aligned}\Delta R_j(q_{ij}) &= p_j q_j - \widetilde{p_j q_j}, \\ \Delta C_i(q_{ij}) &= \theta c_i q_i - \theta \widetilde{c_i q_i}.\end{aligned}$$

Standard derivations lead to the equilibrium price:

$$\begin{aligned}p_{ij} &= (1-\phi) \cdot \frac{\Delta R_j(q_{ij})}{q_{ij}} + \phi \cdot \frac{\Delta C_i(q_{ij})}{q_{ij}} \\ &= (1-\phi) \cdot MC_i(q_{ij}^*) + \phi \cdot AC(q_{ij}^*)\end{aligned}$$

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 the set of feasible equilibria. The efficient quantity is q^*, while the negotiated price ranges between $MC(q^*)$ and $AC(q^*)$ depending on ϕ.

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:

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

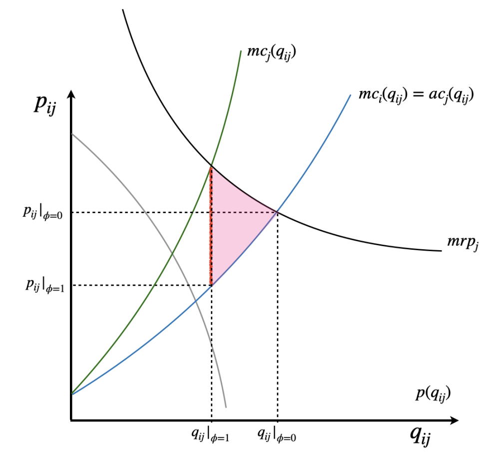

FIGURE B.2: Supply-Driven Quantity Bargaining

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:

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

as in equation (2.1). Solving the bargaining problem using derivations similar to those in Section A.1 of this Appendix yields the following expression for the bilateral price:

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

where

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

Here, ω_{ij}^S is given by:

$$\omega_{ij}^S \equiv \frac{\frac{\phi}{1-\phi} \cdot \lambda_{ij}^S}{1 + \frac{\phi}{1-\phi} \cdot \lambda_{ij}^S} \quad \text{and} \quad \lambda_{ij}^S \equiv \frac{GFT_{ij}^i}{GFT_{ij}^j} \psi_{ij} = \frac{p_{ij} q_{ij} - \theta c_i q_i \Delta_{ij}^x}{p_j q_j \Delta_{ij}^s - p_{ij} q_{ij}} \psi_{ij}.$$

Hence, the bilateral price is a convex combination of the exporter's marginal cost (c_i) and a term reflecting the markdown (ψ_{ij}) below the importer's marginal revenue product ($MRP_j \equiv \frac{dp_j q_j}{dq_{ij}}$). As in the demand-driven quantity bargaining case, the weight ω_{ij}^S has an intuitive interpretation as the importer's effective bargaining power.

For $\phi = 0$, $\omega_{ij}^S = 0$, and the allocation is efficient with $p_{ij} = c_i$. For $\phi = 1$, $\omega_{ij}^S = 1$, and we obtain the standard oligopsony benchmark where the input price is a markdown below the importer's marginal revenue product, with the markdown equal to one plus the importer's residual supply elasticity. This reflects the importer's view, where the cost of increasing quantity is driven by raising the unit price for all inframarginal units when supply curves are upward-sloping.

Figure B.2 plots the set of feasible allocations. The red line indicates the pure oligopsony markdown, and the pink area represents the set of markdowns for $\phi \in (0, 1)$.

While this case is intuitive, it is more complex to characterize due to the term λ_{ij}^S, which is itself a function of prices and quantities and thus no longer a simple function of market shares as in the baseline case. As a result, bringing this model to the data is less straightforward than our baseline model.

B.3 Generalized Outside Option

In the baseline model, we assume that the importer's (exporter's) gains from trade are given by the firm's total payoff from trading with all partners, minus the payoff from trading with all partners except exporter i (importer j). In simpler terms, each importer (exporter) treats itself as the marginal buyer (supplier). This implies that, in case of disagreement, neither party considers the possibility of forming new relationships, which substantially influences the definition of outside profits for both importers and exporters.

We now consider a more general specification that imposes less structure on outside options. Let

$$\Delta_{ij}^{c_i} = \frac{\tilde{c}_i}{c_i} \qquad \text{and} \qquad \Delta_{ij}^{c_j} = \frac{\tilde{c}_j}{c_j}$$

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:

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

where we define:

$$\mu_{ij}^{\text{oligopsony},G} := \theta\left[\frac{\Delta_{ij}^x}{x_{ij}}\right], \qquad \Delta_{ij}^x := \left[\frac{1 - \Delta_{ij}^{c_i}(1 - x_{ij})}{x_{ij}}\right],$$
$$\Delta_{ij}^{s,G} := 1 - \left(\Delta_{ij}^{c_j}\right)^{1-\nu}.$$

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

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

where the weighting factor is given by:

$$\omega_{ij}^G := \frac{\frac{\phi}{1-\phi}\lambda_{ij}^G}{1 + \frac{\phi}{1-\phi}\lambda_{ij}^G} \in (0,1),$$
$$\lambda_{ij}^G := \frac{\gamma s_{ij}}{\Delta_{ij}^{s,G}} \cdot \frac{\nu - 1}{\varepsilon_{ij} - 1} > 0.$$

Equation (B.1) shares the structure of the markup equation in equation (2.6), but with two key differences. First, the term λ_{ij} now depends on $\Delta_{ij}^{s,G}$, which cannot be expressed solely as a function of the supplier share. Second, the oligopsony markdown depends on Δ_{ij}^{x}, which likewise cannot be written as a function of the buyer share alone.

To summarize, generalizing the structure of outside options preserves the overall structure of the markup equation and yields similar comparative statics. However, it introduces an identification problem: the effective bargaining power and oligopsony markdown now depend on unobserved terms, such as $\Delta_{ij}^{s,G}$ and Δ_{ij}^{x}, that cannot be expressed solely as functions of market shares. In the absence of external data to estimate these terms, empirical implementation becomes infeasible without further assumptions. Imposing the more restrictive baseline assumption enables us to express both components as functions of observable bilateral market shares and a small set of structural parameters, allowing us to implement the model using available data and avoid excessive computational complexity.

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 generalize the cost shock to an exporter-level shock, 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:

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

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:

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

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

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

where:

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

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.

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 computed by comparing the NAICS6 industry code listed in ORBIS to the HS6 product code recorded in the customs data, using the concordance developed by Pierce and Schott (2009). We retain only those matches where both location and product scores exceed 90%. In addition, we drop from the matched dataset any ORBIS firm with fewer than five transactions to filter out spurious exporters and potential noise.

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 6% 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's structure and requirements.

We begin by removing observations that are incomplete or inconsistent with the modeling framework. Specifically, we drop transactions with missing or zero values for import value or quantity, invalid exporter identifiers (e.g., strings with fewer than three characters or beginning with a number), or U.S. importers that cannot be linked to the Longitudinal Business Database (LBD). We also exclude transactions associated with special provisions or temporary classifications (HS codes 98–99).

Next, we restrict our attention to trade in intermediate and capital goods by removing HS10 products, which are classified as consumption goods under the BEC system.

Although the customs data are reported at the transaction level, we aggregate them to the annual level for each buyer–supplier–product triplet. We choose annual aggregation because few relationships appear in adjacent months or in the same month across years. We then retain only those triplets that are active in at least two consecutive years and where the supplier transacts the same HS10 product with more than one U.S. buyer. This ensures a panel structure with repeated observations, supporting the identification of the model's parameters.

We also exclude related-party trade, which is less likely to reflect decentralized bargaining and more likely to involve internal pricing practices such as transfer pricing.[35] In our baseline, a buyer–supplier pair is flagged as related if ORBIS identifies a shared corporate parent. To preserve coverage, we retain all observations not flagged as related by either ORBIS or the LFTTD. For robustness, we also consider two alternative definitions: one based solely on the LFTTD flag, and another combining it with ORBIS data to identify the U.S. importer as a multinational. See Appendix C.1 for further details.

To address outliers, we follow Heise (2024) and apply two filters. First, we drop observations with extreme price levels, defined as log unit values below the 1st or above the 99th percentile of the HS10 product–country distribution. Second, we trim extreme price changes by excluding year-on-year log price differences smaller than -4 or greater than $+4$. These outliers are removed from estimation but retained when computing tenure and relationship length. We also exclude all HS10 products under HS chapter 27 (energy-related goods), which lie outside the scope of the model.

[35] Bernard et al. (2006) shows that related-party prices differ systematically from arm's-length transactions, including lower average prices and different pass-through behavior.

After applying these restrictions, which exclude outliers, energy products, and related-party trade, and focusing on repeated, arm's-length relationships, the sample retains over 20% of U.S. imports by value and nearly 15% of buyer–supplier–product triplets.

Finally, we restrict attention to suppliers that sell the same HS10 product to more than one U.S. buyer. This condition is essential for the identification of the model's firm-level parameters and ensures that the empirical setting aligns with the model's structure oulined in Setion 4.

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 show the estimation samples used in the analysis. For 2001–2016 (Panel A), the data include approximately \$880 billion in import value, 480,000 buyer–supplier pairs, and 630,000 buyer–supplier–HS10 triplets. For 2017–2018 (Panel B), the final sample includes \$160 billion in imports, 190,000 pairs, and 250,000 triplets.[36] These samples correspond exactly to the data used in estimation and post-estimation analysis.

Table C.2 reports a similar analysis using a broader sample that includes capital, intermediate, and consumption goods ("BEC – All (Consec.)"), rather than limiting to consumption goods only. Analogous to Table 1, Table C.3 reports the summary statistics of this broader sample.

[36]Panel A omits intermediate sample steps and shows only the full and final samples for 2001–2016, as these are not reported in the analysis nor disclosed under Census data policies. The same applies to Table C.2.

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)
		Panel A: 2001–2016			
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
		Panel B: 2017–2018			
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: Sample Composition by Period – All BEC Categories

Sample	Import value (bn USD)	Importers (th)	Exporters (th)	Pairs (th)	Triplets (th)
Panel A: 2001–2016					
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
Panel B: 2017–2018					
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.3: Summary Statistics – All BEC Categories (2001–2018)

Variable	Mean	Std. Dev.	P25	Median	P75
Panel A: Characteristics of Trade Relationships					
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	—	—	—	—
Panel B: Prices					
$\log p$ (pre-duty)	3.40	2.50	1.50	3.10	5.10
$\log p$ (pre-duty, excl. charges)	3.30	2.50	1.40	3.00	5.00
$\log p^{\text{duty}}$ (post-duty)	3.40	2.50	1.50	3.10	5.20

Notes: This table reports summary statistics for a sample that covers all BEC product categories except energy goods, and excludes statistical outliers and related-party trade. It further restricts to suppliers that trade with at least two different U.S. buyers in consecutive years. This corresponds to the cumulative sample underlying the "+ Supplier Multi-Buyer" row in Panel B of Table C.2. Columns report the mean, standard deviation, and selected quantiles (25th percentile, median, and 75th percentile) for each variable. Prices in Panel B are log unit values (FOB value over quantity), with variants including charges or duties. s_{ijh} denotes exporter i's share in buyer j's imports of product h; x_{ijh} denotes buyer j's share in exporter i's U.S. exports of the same product. Relationship length and tenure are in years; concentration is measured at the HS10–year level. Counts of buyers, suppliers, and origin countries are per product per firm. Statistics are based on confidential LFTTD data and rounded to four significant digits per 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:

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

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

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	$\ln p_{ijht}$	$\ln p^{c}_{ijht}$	$\ln p^{\text{duty}}_{ijht}$
Panel A: Overall price dispersion			
FE_{ht}	0.483	0.485	0.486
FE_{i}	0.427	0.424	0.423
FE_{j}	0.0452	0.0464	0.0463
Match residual	0.0444	0.0441	0.0442
Panel B: Within exporter–product dispersion			
FE_{j}	0.231	0.233	0.233
Match residual	0.768	0.765	0.765

Notes: The columns correspond to alternative price definitions: $\ln p_{ijht}$ excludes both duties and charges; $\ln p^{c}_{ijht}$ includes charges but excludes duties; $\ln p^{\text{duty}}_{ijht}$ includes duties but excludes charges. 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 $\boldsymbol{X}_{ijht}$ 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 1,2000,000 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 the heterogeneity in tariff pass-through through specification (3.2) by including buyer-by-year fixed effects (FE_{jt}). This more demanding specification accounts for time-varying shocks at the buyer level, while also controlling for product–year and exporter country–year fixed effects ($FE_{ht} + FE_{ct} + FE_{jt}$). Findings replicate the results of incomplete tariff pass-through reported in Table 3, suggesting that exporters adjust marginal costs in response to demand shifts from dominant importers, thereby absorbing a substantial fraction of tariff shocks. These results demonstrate the critical role of the cost channel as, by Proposition 3, the pass-through decreases with the buyer share x_{ij}.

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

Dependent variable:	$\Delta \ln p_{ijht}$					
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1 + \tau_{cht})$	-0.223	-0.342	-0.230	-0.154	-0.163	-0.292
	(0.109)	(0.155)	(0.096)	(0.130)	(0.107)	(0.145)
$\Delta \ln(1 + \tau_{cht}) \cdot \ln \text{longevity}_{ijht}$		0.086				0.097
		(0.044)				(0.044)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$			0.023		0.029	0.018
			(0.141)		(0.157)	(0.157)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$				-0.271	-0.271	-0.280
				(0.127)	(0.135)	(0.130)
$FE_{ht} + FE_{ct} + FE_{jt}$	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 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.

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

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 in magnitude and statistically insignificant, indicating that supplier concentration does not materially affect the degree of pass-through. In contrast, buyer shares exhibit a strong, monotonic relationship (in green): pass-through declines significantly at higher quartiles of x_{ijht}, consistent with the model's prediction that dominant buyers constrain suppliers' ability to shift cost shocks.

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

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

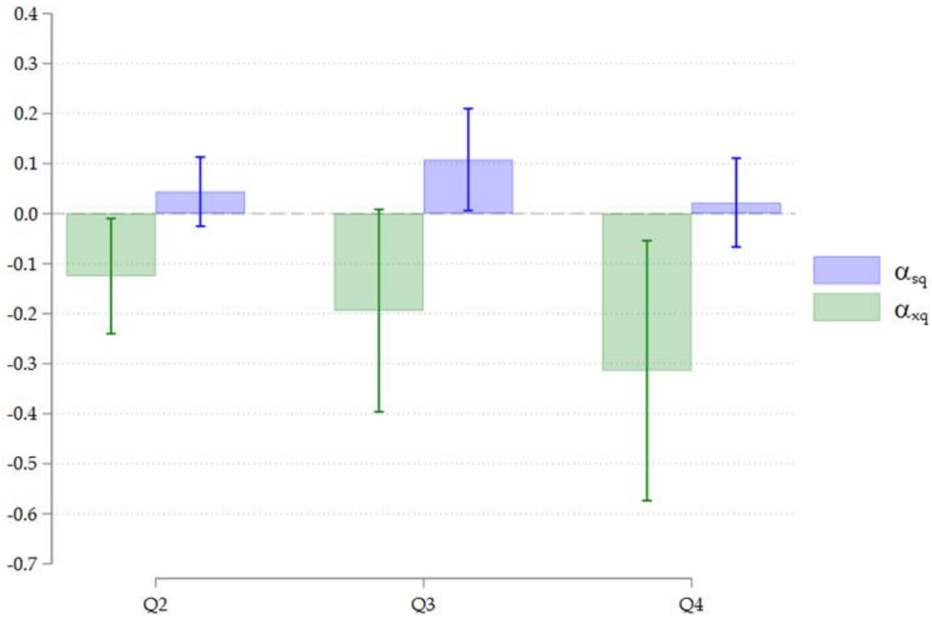

(B) Fixed Effects: $FE_{ht} + FE_{ct} + FE_{jt}$

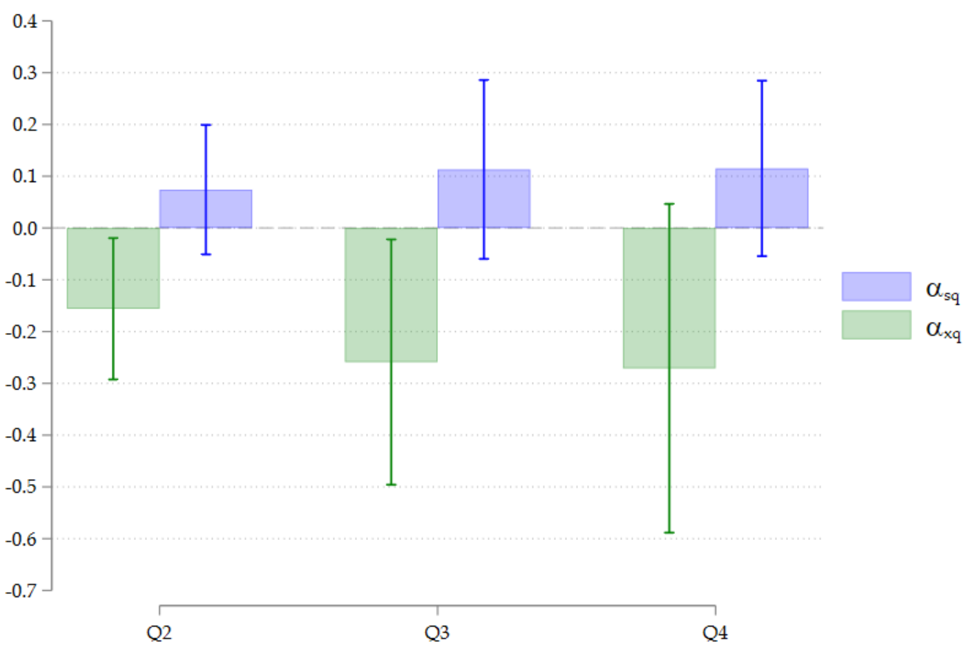

Notes: The figure plots estimated coefficients from regressions of bilateral price changes on tariff changes interacted with quartiles of supplier share (s_{ijht}) and buyer share (x_{ijht}). The first quartile serves as the omitted category. The estimated coefficients correspond to equation (D.1), where $\alpha_{s,q}$ and $\alpha_{x,q}$ capture the interaction of the tariff term with the qth 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 from our baseline model. Columns (3)-(4) use duty-exclusive prices that *include* charges, while Columns (5)-(6) use tariff-inclusive prices. In all cases, we interact tariff changes with lagged supplier and buyer shares, and hold fixed effects constant across specifications for comparability. Across all variations, pass-through estimates and their interaction effects with bilateral market shares remain stable in sign and magnitude, supporting the robustness of the main findings.

TABLE D.3: Additional Robustness: GE Controls and Price Definitions

Dependent variable:	$\ln p_{ijht}$ (excl. GE controls)		$\ln p^{c}_{ijht}$ (before duty, incl. charges)		$\ln p^{\text{duty}}_{ijht}$ (tariff inclusive)	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1 + \tau_{cht})$	-0.060	-0.153	-0.032	-0.114	0.506	0.401
	(0.093)	(0.106)	(0.091)	(0.104)	(0.099)	(0.126)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$	0.048	0.036	0.033	0.004	0.017	0.007
	(0.069)	(0.162)	(0.080)	(0.165)	(0.068)	(0.143)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$	-0.399	-0.278	-0.407	-0.280	-0.421	-0.287
	(0.113)	(0.136)	(0.112)	(0.134)	(0.121)	(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 including 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 replicate the main GMM estimation using an expanded sample that includes all Broad Economic Categories (BECs), rather than excluding consumption goods as in the baseline. This extended product scope allows us to test whether our key parameter estimates, namely, returns to scale (θ) and relative bargaining power (ϕ), are sensitive to the exclusion of consumer-oriented products. As shown in Table D.4, the results remain fairly stable, suggesting that the baseline findings are not driven by product composition.

We then examine how sensitive the estimates are to alternative values of calibrated model parameters. First, we vary the elasticity of substitution across foreign varieties (ρ) by setting $\rho = 5$ instead of 10, which is consistent with the lower end of the estimates in the literature. Second, we relax the assumption of constant returns to scale for the importers and set $\varrho = 0.5$ instead of 1. We report the results in Table D.5. Column (1) shows the estimated values when setting $\rho = 5$, and Column (2) shows the estimated values when setting $\varrho = 0.5$. Furthermore, in Table D.5, we also estimate the parameters using an alternative sample constructed by utilizing the RPT indicator from LFTTD.[37] Columns (3) and (4) report the estimated values for this set of sample. Throughout these alternative setups,the resulting estimates remain robust, suggesting that the estimated values are not sensitive to particular values of other parameters or set of sample.

[37] See Section 3.3 for the discussion on selection through RPT indicators.

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: $\widehat{\ln \frac{\phi}{1-\phi}}$	1.162		0.558	
	(0.026)		(0.022)	
Returns to scale ($\hat{\theta}$)	0.505	0.573	0.432	0.586
	(0.003)	(0.005)	(0.004)	(0.005)
Constant		2.780		0.742
		(0.190)		(0.083)
Longevity		0.061		0.742
		(0.023)		(0.069)
Number of HS10 transactions		-0.240		-0.019
		(0.019)		(0.011)
Multiple HS10 dummy		0.080		0.190
		(0.026)		(0.025)
Lagged outside option		-0.188		-0.237
		(0.017)		(0.021)
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	0.762	0.896	0.636	0.848
	(0.005)	(0.068)	(0.005)	(0.112)
Median	–	0.913	–	0.878
	–	(0.068)	–	(0.112)

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: the elasticity of demand (ν), the elasticity of costs with respect to foreign input prices (γ), and the elasticity of substitution across foreign varieties (ρ). We set $\varrho = 1$, so that $\eta = 2.5$. Panel B presents GMM estimates. Columns (1) and (3) impose a constant ϕ across bilateral pairs, while Columns (2) and (4) estimate the full vector $\boldsymbol{\kappa}$ to allow for heterogeneity in bargaining power. Specifications differ in the inclusion of fixed effects. Controls include: (i) the log of relationship longevity between exporter i and importer j; (ii) the log of the number of transactions between i and j in a given year; (iii) the log of the relative outside option, defined as the ratio of exporter i's sales to other U.S. buyers (excluding j) over importer j's purchases from other suppliers (excluding i), both in year $t-1$; and (iv) a dummy variable equal to one if the i–j pair transacts in more than one HS10 product. Panel C reports the mean and median of the implied bargaining power. Standard errors are robust; those in Panel C are computed using the delta method. The set of instruments includes the number of exporters and importers at the HS10 level, as well as lagged bilateral shares (excluding the focal pair). 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: Robustness of Model Estimates

	$\rho = 5$	$\varrho = 0.5$	Related Party Trade LFTTD	
	(1)	(2)	(3)	(4)
Rel. bargaining power: $\widehat{\ln \frac{\phi}{1-\phi}}$	1.455	1.838	1.618	0.892
	(0.038)	(0.052)	(0.062)	(0.047)
Returns to scale (θ)	0.427	0.473	0.453	0.381
	(0.004)	(0.003)	(0.005)	(0.007)
Mean $\hat{\phi}$	0.811	0.863	0.835	0.709
	(0.006)	(0.006)	(0.009)	(0.010)
Median $\hat{\phi}$	0.811	0.863	0.835	0.709
	(0.000)	(0.000)	(0.000)	(0.000)
None	Yes	Yes	Yes	No
$FE_h + FE_t + FE_j$	No	No	No	Yes
Observations		3,120,000		

Notes: This table reports robustness checks for the main model estimates. The columns explore sensitivity to changes in key calibrated parameters and sample definitions. Column (1) varies the elasticity of substitution across foreign varieties (ρ); Column (2) changes the downstream returns to scale parameter (ϱ). Columns (3)–(4) use a sample based on related-party indicators from LFTTD (RPT), without and with fixed effects for buyer, product, and time. 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

Sensitivity to alternative parameter values. Table D.6 presents IV-based goodness-of-fit tests as in equation (4.5) for an 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 3.

Quantity responses and relationship heterogeneity: data vs. model. Table D.7 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.

Testing the model-predicted quantity changes. Panel A in Table D.8 presents IV-based tests comparing observed quantity changes to those predicted by the model under alternative parameterizations. The model-predicted change in quantity is given by:

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

where ε_{ijht} is the match-specific demand elasticity and $\widehat{\Delta \ln q_{ijht}}$ is as in equation (4.5). Although formal tests reject all models, the baseline model with bargaining and decreasing returns (Columns (1) and (2)) demonstrates the strongest fit, indicating that this specification best captures the underlying mechanisms of tariff-induced quantity adjustments.

Testing the model-predicted sales changes. Panel B in Table D.8 provides analogous IV-based tests for observed sales changes computed as:

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

The baseline specification (Columns (1) and (2)) shows superior alignment between predicted and observed sales, reinforcing the conclusion that the bargaining model with decreasing returns most effectively matches empirical patterns in the data.

TABLE D.6: IV-Based Goodness-of-Fit Test, Alternative Parameters

Panel A: Baseline Fixed Effects

Dependent variable:	$\Delta \ln p_{ijht}$							
	$\rho = 5$				$\varrho = 0.5$			
	Baseline	$\phi = 0$, $\theta = 1$	$\theta = 1$	$\phi = 0$	Baseline	$\phi = 0$, $\theta = 1$	$\theta = 1$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \widehat{\ln p_{ijht}}$	1.211	0.781	0.735	1.368	1.247	0.871	0.740	1.544
	(0.343)	(0.220)	(0.207)	(0.387)	(0.352)	(0.246)	(0.209)	(0.437)
$FE_{ht} + FE_{cs}$	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Stringent Fixed Effects

Dependent variable:	$\Delta \ln p_{ijht}$							
	Baseline	$\phi = 0$, $\theta = 1$	$\theta = 1$	$\phi = 0$	Baseline	$\phi = 0$, $\theta = 1$	$\theta = 1$	$\phi = 0$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \widehat{\ln p_{ijht}}$	1.000	0.648	0.612	1.132	1.047	0.720	0.615	1.302
	(0.351)	(0.227)	(0.214)	(0.398)	(0.368)	(0.252)	(0.216)	(0.458)
$FE_{ht} + FE_{ct} + FE_{jt}$	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 price on the corresponding model-predicted change $\Delta \widehat{\ln p_{ijht}}$, 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 $\varrho = 0.5$ instead of $\varrho = 1$ 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.7: *Quantity* Responses and Relationship Heterogeneity: Data vs. Model

	(1)	(2)	(3)	(4)
Panel A: Data				
$\Delta \ln(1+\tau_{cht})$	-0.568	0.006	-1.021	-0.761
	(0.249)	(0.331)	(0.185)	(0.237)
$\Delta \ln(1+\tau_{cht}) \cdot s_{ijht-1}$		-0.511		-0.278
		(0.323)		(0.276)
$\Delta \ln(1+\tau_{cht}) \cdot x_{ijht-1}$		-0.201		-0.130
		(0.229)		(0.242)
R-squared	0.06	0.10	0.32	0.36
Panel B: Model				
	(1)	(2)	(3)	(4)
$\Delta \ln(1+\tau_{cht})$	-1.962	-3.967	-2.169	-3.984
	(0.080)	(0.201)	(0.108)	(0.230)
$\Delta \ln(1+\tau_{cht}) \cdot s_{ijht-1}$		2.441		2.420
		(0.409)		(0.551)
$\Delta \ln(1+\tau_{cht}) \cdot x_{ijht-1}$		2.993		3.270
		(0.229)		(0.316)
R-squared	0.28	0.38	0.44	0.51
$FE_{ht} + FE_{cs}$	Yes	Yes	No	No
$FE_{ht} + FE_{ct} + FE_{jt}$	No	No	Yes	Yes
Observations	249,000			

Notes: This table reports the pass-through of tariffs to quantities at the exporter–importer–product level. Panel A presents reduced-form estimates from the data. Panel B shows corresponding pass-through estimates generated by the model. Columns (2)–(4) interact tariff changes with 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) employ a more stringent 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.8: IV-Based Goodness-of-Fit Test for Quantities and Sales

Panel A: Quantities

Dependent Variable:	$\Delta \ln q_{ijht}$							
	Baseline		$\phi = 0,\ \theta = 1$		$\theta = 1$		$\phi = 0$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta \widehat{\ln q_{ijht}}$	0.349	0.525	0.205	0.311	0.185	0.282	0.420	0.627
	(0.144)	(0.142)	(0.084)	(0.084)	(0.076)	(0.076)	(0.174)	(0.170)

Panel B: Sales

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)
$\Delta \widehat{\ln(p \cdot q)}_{ijht}$	0.176	0.426	0.102	0.250	0.092	0.228	0.212	0.509
	(0.157)	(0.154)	(0.091)	(0.090)	(0.082)	(0.082)	(0.189)	(0.184)
$FE_{ht} + FE_{cs}$	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

Panel A and Panel B in Table D.9 evaluate how tariff pass-through estimates vary under different sample restrictions using baseline fixed effects (product–time, country–sector) and more stringent fixed effects (including buyer-time), respectively. For ease of exposition, the first two columns of Table D.9 replicate the numbers of Columns (1) and (3) of Table 5. These two columns show that the baseline sample yields a pass-through estimate of around 78-85%, closely matched by the model-implied prediction with no statistically significant difference. As the sample is progressively broadened–from including single-buyer matches (Column (3)), to adding related-party transactions, energy goods, and outliers (Column (4)), and finally to the most inclusive specification (Column (5))–estimated pass-through increases steadily, reaching up to 93-95%. This pattern underscores the sensitivity of reduced-form estimates to sample composition and the role of relationship filtering in uncovering pricing patterns consistent with bilateral bargaining.

Table D.9: Tariff Pass-Through Across Different Samples

Dependent variable: $\Delta \ln p_{ijht}$	Model	Baseline	+Suppliers w/ < 2 Buyers	+ Energy/ RPT/ Outliers	+ Final Goods
	(1)	(2)	(3)	(4)	(5)
Panel A: Baseline Fixed Effects					
$\Delta \ln(1+\tau_{cht})$	-0.248	-0.151	-0.168	-0.099	-0.066
	(0.008)	(0.093)	(0.035)	(0.050)	(0.043)
$FE_{ht} + FE_{cs}$	Yes	Yes	Yes	Yes	Yes
R-squared	0.32	0.04	0.02	0.02	0.02
Panel B: Stringent Fixed Effects					
$\Delta \ln(1+\tau_{cht})$	-0.249	-0.223	-0.171	-0.123	-0.047
	(0.010)	(0.109)	(0.039)	(0.059)	(0.042)
$FE_{ht} + FE_{ct} + FE_{jt}$	Yes	Yes	Yes	Yes	Yes
R-squared	0.50	0.31	0.18	0.14	0.13
Observations	249,000	249,000	732,000	945,000	1,768,000

Notes: This table reports tariff pass-through estimates to duty-exclusive 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).

D.5.2 Pass-Through Using Data Aggregated at the Product-Level

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–country–month level, using the same subset of buyer–supplier–product links as in the baseline analysis. These are links with at least one transaction in both 2017 and 2018. For each product and country, we construct monthly prices by aggregating trade values and quantities. We then relate monthly price changes to changes in statutory tariffs, controlling for product-month, country-month, and country-sector fixed effects. Standard errors are clustered at the HS8–country level.

Models Table D.10, reports results separately for all products and for the subset that excludes consumption goods. Within each group, columns reflect increasingly selective samples. Columns (1) and (4) includes buyer–supplier–product pairs observed in two consecutive calendar years. The next specification, Columns (2) and (5), excludes consumption goods and applies additional filters: energy-sector goods are dropped; transactions with extreme price levels and price changes are excluded; and related-party trade is removed. The final sample, Columns (3) and (6), is restricted to suppliers trading with at least two U.S. buyers for the same product in consecutive years.

Results Across specifications, we find consistent evidence of incomplete tariff pass-through to U.S. import prices. As in our firm-level regressions, the degree of pass-through incompleteness increases as we move to more selective samples, particularly those that condition on firms with multiple trading partners over consecutive years. The estimated effects are generally larger in magnitude than at the match level, suggesting that relationship-level frictions may be amplified when observed in aggregated trade flows. While this exercise remains suggestive, it helps connect the mechanisms explored in the main analysis to pricing patterns in product-level data.

Table D.10: Tariff Pass-Through Using Data Aggregated at the Product-Level

Dependent Variable:	$\Delta \ln p_{cht}$					
	All Products			Excl. Consumption Goods		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \ln(1+\tau_{cht})$	-0.115	-0.136	-0.321	-0.072	-0.119	-0.700
	(0.078)	(0.058)	(0.186)	(0.136)	(0.101)	(0.164)
Consecutive obs only	Yes	Yes	Yes	Yes	Yes	Yes
(-) Energy/RPT/Outliers	No	Yes	Yes	No	Yes	Yes
(-) Suppliers with <2 buyers	No	No	Yes	No	No	Yes
Observations	800,000	540,000	180,000	510,000	320,000	100,000
R-squared	0.13	0.14	0.22	0.13	0.14	0.22

Notes: This table reports regressions of month-over-month changes in log unit values (FOB, excluding charges) on corresponding changes in statutory tariffs ($\Delta\tau_{hct}$), measured at the HS10–country–month level. The data are aggregated from the firm-level sample used in our main analysis, retaining only buyer–supplier–product triplets with consecutive transactions. Columns (1)-(3) refer to all products; Columns (4)-(6) exclude consumption goods, based on the Broad Economic Categories (BEC) classification. Each column reflects a progressively more restricted sample: Columns (1) and (4) include all consecutive transactions; Columns (2) and (5) drop energy goods, extreme price levels and changes, and related-party trade; and Columns (3) and (6) restrict to suppliers with multiple buyers in consecutive years. All regressions include product–month, country–month, and country–sector fixed effects. Standard errors are clustered at the HS8–country level and reported in brackets. 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.5.3 Pass-Through Results Using LFTTD RPT Indicators

As a robustness check, we re-estimate the baseline specifications using the related-party transaction (RPT) indicator provided in the LFTTD. Table D.11 shows that the pass-through estimates remain stable, and the interaction effects with supplier and buyer shares are qualitatively similar to those reported in Table 3.

Table D.11: Pass-Through and Relationship Heterogeneity–Alternative Related Party Trade

Dependent variable:	$\Delta \ln p_{ijht}$			
	(1)	(2)	(3)	(4)
$\Delta \ln(1 + \tau_{cht})$	-0.163	-0.087	-0.304	-0.253
	(0.096)	(0.100)	(0.111)	(0.114)
$\Delta \ln(1 + \tau_{cht}) \cdot s_{ijht-1}$		0.043		0.011
		(0.070)		(0.174)
$\Delta \ln(1 + \tau_{cht}) \cdot x_{ijht-1}$		-0.374		-0.222
		(0.110)		(0.148)
$FE_{ht} + FE_{cs}$	Yes	Yes	No	No
$FE_{ht} + FE_{ct} + FE_{jt}$	No	No	Yes	Yes
Observations	249,000	249,000	249,000	249,000
R-squared	0.04	0.05	0.32	0.32

Notes: This table reports estimates of the pass-through of statutory tariffs, $\Delta \ln(1 + \tau_{cht})$, to duty-exclusive prices at the exporter–importer–product–year level, $\Delta \ln p_{ijht}$. Columns (2) and (4) include interactions between tariffs and lagged bilateral characteristics: supplier share (s_{ijht-1}), defined as supplier i's share in buyer j's imports of product h, and buyer share (x_{ijht-1}), defined as buyer j's share in supplier i's exports of product h. All regressions include the following controls: (i) $\ln \text{longevity}_{ijht}$, the number of years i and j have transacted in h; (ii) $\Delta \ln q_{i(-j)ht}$, exporter i's sales of h to U.S. buyers other than j; and (iii) $\Delta \ln p_{(-i)jht}$, the average price change charged by other suppliers of h to buyer j, using lagged shares as weights. Columns (1)–(2) include product–year and exporter country–sector fixed effects ($FE_{ht} + FE_{cs}$), while Columns (3)–(4) include a more demanding set of fixed effects: product–year, importer–year, and exporter country–year ($FE_{ht} + FE_{ct} + FE_{jt}$). Standard errors are clustered at the HS8 product and exporter-country level. Arm's length transactions are defined using LFTTD related party trade indicator. 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).

E Estimation Appendix

E.1 Downstream Demand Elasticity (ν)

Consider a model where importer j sells its output q_j to downstream customers in different countries. A representative consumer in each country maximises utility by choosing a composite of domestic and imported goods. The sub-utility derived from the composite imported good will be given by a CES aggregation across imported varieties with a good-importer specific elasticity of substitution given by σ_g. Broda and Weinstein (2006) provide estimates of the elasticity σ_g at the HS10 good g-level in U.S. import data. The plot below shows the distribution of these elasticities. We base the calibration of the elasticity ν in our model on these estimates. We consider a value of 4 for ν, close to the mean value of 3.85, which we see as a conservative choice.

FIGURE E.1: Downstream Demand Elasticity

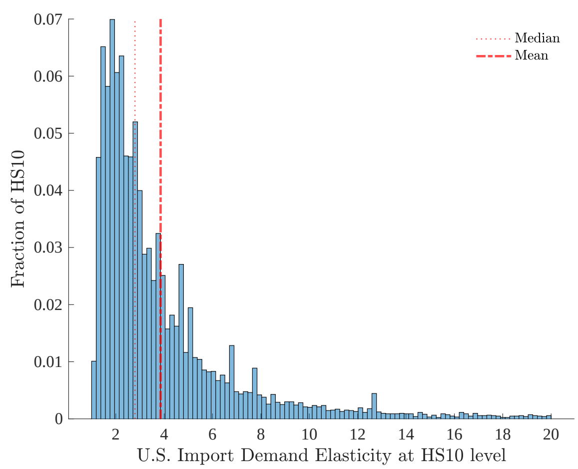

Notes: The figure displays the estimates of the import demand elasticity σ_g from Broda and Weinstein (2006). The mean and median value of σ_g^{US} is 3.85 and 2.8, respectively. Estimates are truncated above at 20, and below at 1.

E.2 Monte Carlo Simulation

Data for one replicate. Each exporter $i \in \{1, \ldots, 200\}$ belongs to a block with exactly two importers, labeled $j(i)$ and $\ell(i)$. Store the log-price difference

$$\Delta p_i = \ln p_{ij} - \ln p_{i\ell} \quad \text{and the pair } (j, \ell).$$

We set the marginal cost to 1 for all pairs for simplicity and use the parameters of $\varrho = 1, \nu = 4, \gamma = 0.5, \rho = 10$. $\phi^* = 0.827$ and $\theta^* = 0.454$. s_{ij} and x_{ij} are drawn from a $U[0,1]$ so that all shares within a block sum to 1. All Monte-Carlo exercises use 501 random replicas.

Joint estimation of ϕ and θ.

1. *Candidate markups.* For any $(\phi, \theta) \in (0,1) \times (0,1)$ compute the bilateral markup $\mu_{ij}(\phi, \theta)$ from the structural formula (2.6).

2. *Model-implied gap for exporter i:* $\Delta\mu_i(\phi, \theta) \; = \; \ln \mu_{ij}(\phi, \theta) \; - \; \ln \mu_{i\ell}(\phi, \theta)$.

3. *Non-linear least squares criterion.*

$$Q(\phi, \theta) \; = \; \sum_{i=1}^{200} \Big[\, \Delta p_i - \Delta\mu_i(\phi, \theta) \Big]^2 .$$

4. *Estimation.* Minimize $Q(\phi, \theta)$ subject to the simple box constraints

$$0.01 \leq \phi \leq 0.99, \qquad 0.01 \leq \theta \leq 1.$$

   We record the resulting estimates $\hat{\phi}, \hat{\theta}$ for each of the 500 Monte-Carlo replicates.

Results. Figure E.2 reports the frequencies for the jointly estimated (ϕ, θ) under the parameters noted above. The left panel refers to $\hat{\phi}$ and the right panel to $\hat{\theta}$. Similarly, Figure E.3 shows the analogous results for an alternative set of parameters of $\varrho = 1, \nu = 2.5, \gamma = 0.5, \rho = 5$. Across the figures, the true and average estimates are numerically close, and the distributions of the estimates are centered around the true parameters, showing that our estimator is *consistent*.

E.3 Estimation Under $\theta = 1$

In Section E.3.1, we first demonstrate that assuming $\theta = 1$ in the estimation leads to overestimating ϕ when the true parameters are $(\phi^*, \theta^*) \in (0,1)^2$. Further, in Section E.3.2, we show that using ϕ^* (instead of the overestimated ϕ) in the model validation exercise in Section 4.3 where θ is imposed to be 1 would yield an upper bound on the attainable correlation between model-predicted and observed price changes.

FIGURE E.2: Estimated (ϕ, θ) when $\varrho = 1, \nu = 4, \gamma = 0.5, \rho = 10$

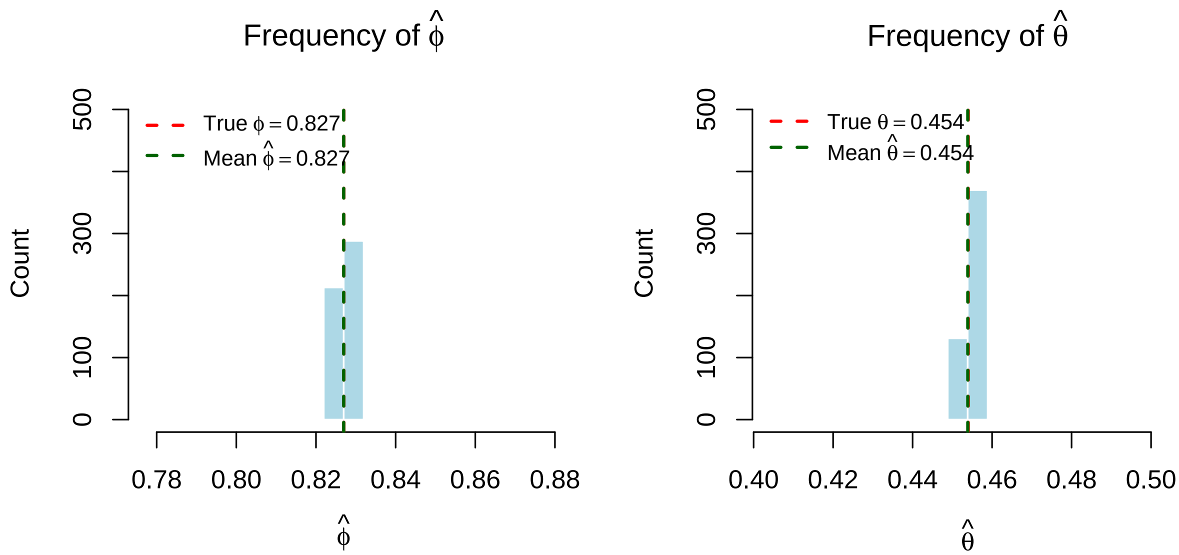

E.3.1 Estimation Bias When Fixing $\theta = 1$ in the Estimation

In what follows, we maintain the following simplification assumptions. We assume $Cor(s_{ij}, x_{ij}) = 0$, mirroring the low correlation between the two bilateral shares (Table 1). We also impose marginal costs to be constant across firms, implying $p_{ij} = \mu_{ij}$ (can be relaxed by assuming a distribution for $\Delta k_{ij\ell}$). Finally, we assume that $\omega(\phi) = \phi$ to keep the notation clean. This assumption is without loss since ω is increasing in ϕ.

The objective function when jointly estimating ϕ and θ is:

$$\begin{aligned} \arg\min_{\phi,\theta} R(\phi, \theta) &= \mathbb{E}\big[p_{ij} - p_{i\ell} - \mu_{ij}(\phi, \theta) + \mu_{i\ell}(\phi, \theta)\big]^2 \\ &= \mathbb{E}\big[\Delta p_i - \Delta \mu_i(\phi, \theta)\big]^2 \end{aligned} \tag{E.1}$$

where j and ℓ are two importers to exporter i and we dropped the ln sign to simplify the notation. Under the full-rank condition, minimization of (E.1) leads to the estimation of the true parameters (θ^*, ϕ^*).

Now define $\tilde{R}(\phi) = R(\phi, \theta)|_{\theta=1}$. When fixing $\theta = 1$, the objective function (E.1) becomes:

$$\arg\min_{\phi} \tilde{R}(\phi) = \mathbb{E}\big[\Delta p_i - \Delta \mu_i(\phi; \theta = 1)\big]^2. \tag{E.2}$$

To study the bias, we can replace $p_{ij} = (1 - \phi^*)\mu_{ij}^{oligopoly} + \phi^* \mu_{ij}^{oligopsony}(\theta^*)$ and rearrange

FIGURE E.3: Estimated (ϕ, θ) when $\varrho = 1, \nu = 2.5, \gamma = 0.5, \rho = 5$

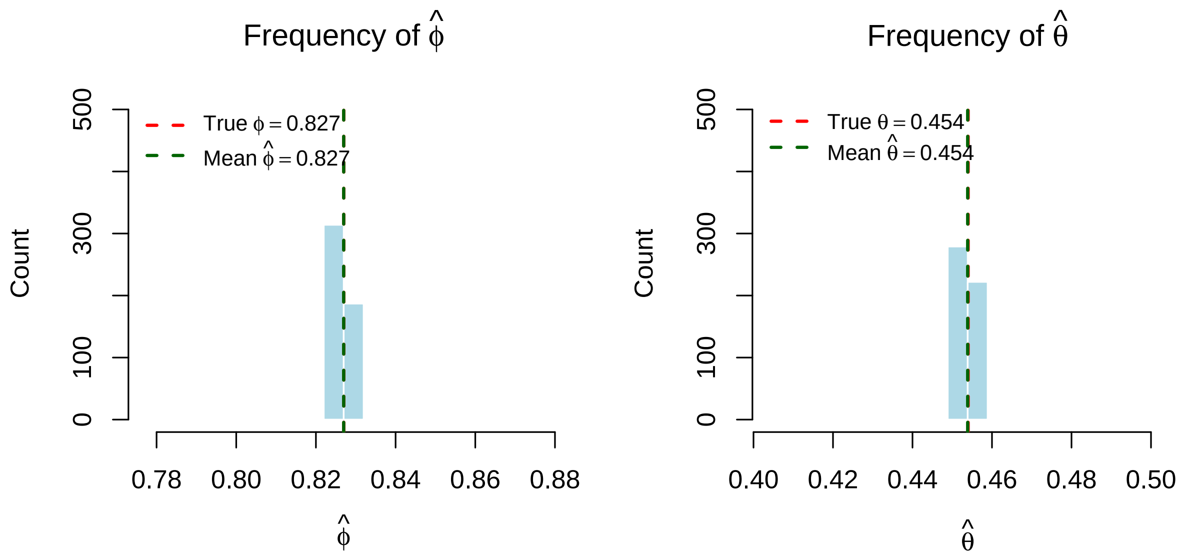

equation (E.2) as follows:

$$\tilde{R}(\phi) = \mathbb{E}\left[(\phi - \phi^*)\left(\mu_{ij}^{oligopoly} - \mu_{i\ell}^{oligopoly}\right) + \phi^*\left(\mu_{ij}^{oligopsony}(\theta^*) - \mu_{i\ell}^{oligopsony}(\theta^*)\right)\right]^2,$$

where, in a slight abuse of notation, we denote the oligopsony markdown computed at the true θ^* as $\mu_{ij}^{oligopsony}(\theta^*)$ to distinguish it from the oligopsony markdown when $\theta = 1$, which is $\mu_{ij}^{oligopsony}(\theta = 1) = 1$.

We can use the price equation to transform this into a function of only prices (data) and $\mu^{oligopsony}(\theta^*)$:

$$\begin{aligned}
\tilde{R}(\phi) =& \mathbb{E}\Bigg[(\phi - \phi^*)\left(\frac{p_{ij} - \phi^*\mu_{ij}^{oligopsony}(\theta^*)}{1 - \phi^*} - \frac{p_{i\ell} - \phi^*\mu_{i\ell}^{oligopsony}(\theta^*)}{1 - \phi^*}\right) \\
& \qquad + \phi^*\left(\mu_{ij}^{oligopsony}(\theta^*) - \mu_{i\ell}^{oligopsony}(\theta^*)\right)\Bigg]^2 \\
=& \mathbb{E}\Bigg[\frac{\phi - \phi^*}{1 - \phi^*}(p_{ij} - p_{i\ell}) - \frac{\phi - \phi^*}{1 - \phi^*}\phi^*\left(\mu_{ij}^{oligopsony}(\theta^*) - \mu_{i\ell}^{oligopsony}(\theta^*)\right) \\
& \qquad + \phi^*\left(\mu_{ij}^{oligopsony}(\theta^*) - \mu_{i\ell}^{oligopsony}(\theta^*)\right)\Bigg]^2 \\
=& \frac{1}{(1 - \phi^*)^2}\mathbb{E}\left[(\phi - \phi^*)(p_{ij} - p_{i\ell}) + \left(\mu_{ij}^{oligopsony}(\theta^*) - \mu_{i\ell}^{oligopsony}(\theta^*)\right)\left((1 - \phi)\phi^*\right)\right]^2
\end{aligned}$$

$$=\frac{1}{(1-\phi^*)^2}\mathbb{E}\left[\left(\phi-\phi^*\right)\Delta p_i+\Delta\mu_i^{oligopsony}\Big(\left(1-\phi\right)\phi^*\Big)\right]^2.$$

Notice that when $\phi=\phi^*$, the first term is zero, but the second term is not. The last equation above can be further rewritten as

$$\begin{aligned}\tilde{R}(\phi)=&\left(\phi-\phi^*\right)^2\cdot\left(\sigma_p^2+\overline{\Delta p}^2\right)+\left(\left(1-\phi\right)\cdot\phi^*\right)^2\cdot\left(\sigma^2_{\Delta\mu^{oligopsony}}+\overline{\Delta\mu}^2_{oligopsony}\right)\\&+2\left(\phi-\phi^*\right)\cdot\left(1-\phi\right)\cdot\phi^*\cdot\left(\mathrm{Corr}\left(\Delta p_i,\Delta\mu_i^{oligopsony}\right)\cdot\sigma_p\cdot\sigma_{\Delta\mu^{oligopsony}}+\overline{\Delta p}\cdot\overline{\Delta\mu}_{oligopsony}\right),\end{aligned}\tag{E.3}$$

where we denote the standard deviation and average of variable z by σ_z and $\overline{z}$ respectively.

Taking derivative with respect to ϕ and setting it equal to zero yields:

$$\phi=\frac{A\cdot\phi^*+B\cdot(\phi^*)^2-\phi^*\cdot C\cdot(1+\phi^*)}{A+B\cdot(\phi^*)^2-2\cdot\phi^*C}=\phi^*\cdot\underbrace{\frac{A+B\cdot\phi^*-C\cdot(1+\phi^*)}{A+B\cdot(\phi^*)^2-2\cdot\phi^*C}}_{>1}>\phi^*$$

where $A=\sigma_p^2+\overline{\Delta p}^2>0,\ B=\sigma^2_{\mu^{oligopsony}}+\overline{\Delta\mu}^2_{oligopsony}>0,\ C=\mathrm{Corr}\left(\Delta p_i,\Delta\mu_i^{oligopsony}\right)\cdot\sigma_p\cdot\sigma_{\mu^{oligopsony}}+\overline{\Delta p}\cdot\overline{\Delta\mu}_{oligopsony}>0$.

Therefore, the objective (E.2) would estimate $\phi>\phi^*$ when setting $\theta=1$. This argument can be extended to heterogeneous marginal costs across pairs. In that case, we would replace $\Delta\mu_i$ for Δp_i.

Simulation We repeat the simulations from Appendix E.2, now estimating only ϕ (with θ fixed at 1). We draw $s_{ij}\sim U[0,1]$, injecting enough within-exporter variation to identify ϕ from differences in markups. If s_{ij} varies very little for exporter i, then $\ln\mu(s_{ij},\phi)-\ln\mu(s_{i\ell},\phi)\approx 0$, so ϕ cannot be identified as the markup difference cancels out in the moment (4.1).

Panel (A) of Figure E.4 maintains 2 importers and exporters in each market, showing a large bias. In Panel (B) of Figure E.4, we increase the number of buyers and suppliers to create more cross-sectional variation. The figure confirms our argument above: The estimated value of ϕ is larger than the true ϕ^* when the estimation imposes $\theta=1$.

E.3.2 Correlation between Pass-Through and Price Changes

In the previous section we proved that the estimated ϕ is necessarily larger than the true value of ϕ^*, when one imposes $\theta=1$ in the estimation. In the exercise of Section 4.3, we do

FIGURE E.4: Estimation of the parameter ϕ while fixing $\theta = 1$

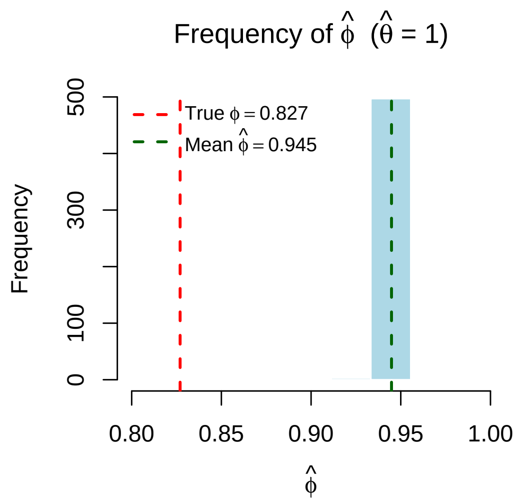

(A) 2 importers and exporters

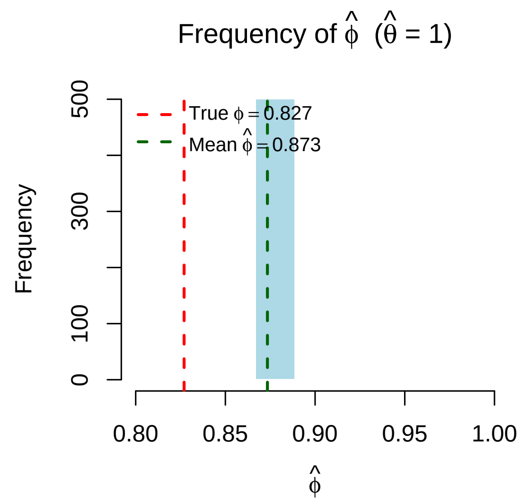

(B) 10 importers and exporters

Notes: In each similuation there are either 2 (Panel a) or 10 (Panel b) buyers and suppliers in each market. There are 100 markets. 501 simulations. We set $\varrho = 1, \nu = 4, \gamma = 0.5, \rho = 10$ and compute ϕ by minimum distance.

not re-estimate ϕ and use ϕ^* when testing the model under $\theta = 1$.

In this section, we show that the coefficients reported for $\theta - 1$ of Figure 3 are *upper bounds* on the attainable correlations between predicted and observed price changes.

In what follows, we drop ln and subscripts to simplify the notation. Call the re-estimated ϕ when setting $\theta = 1$ as ϕ^R. Then, combining equations (4.5) and (4.4) in Section 4.3, the pass-through regression for this model is

$$\Delta p \;=\; \beta^R \cdot \Phi^R \cdot \Delta T \;+\; u^R, \tag{E.4}$$

where we denoted $\Phi^R = \Phi(s, x; \phi^R, \theta = 1)$.[38] We also assume that the residual is independent from $\Phi^R \cdot \Delta T$.

From Appendix E.3.1, setting $\theta = 1$ and re-estimating ϕ, we would find $\phi^R > \phi^*$, so that

$$\tilde{\Phi} \;=\; \Phi(s, x; \phi^*, 1) \;<\; \Phi(s, x; \phi^R, 1) \;=\; \Phi^R,$$

since the pass-through increases in ϕ.

[38]We disregard the fixed effects and consider each variable as demeaned for simplicity.

Using $\tilde{\Phi}$ as an independent variable instead of Φ^R in the OLS regression (E.4),

$$\Delta p = \tilde{\beta} \cdot \tilde{\Phi} \cdot \Delta T + \tilde{u}$$

yields the following estimated coefficient:

$$\tilde{\beta} = \frac{\mathrm{Cov}(\tilde{\Phi}\,\Delta T,\ \Delta p)}{\mathrm{Var}(\tilde{\Phi}\,\Delta T)} = \beta^R\,\frac{\mathrm{Cov}(\tilde{\Phi}\,\Delta T,\ \Phi^R\,\Delta T + u^R)}{\mathrm{Var}(\tilde{\Phi}\,\Delta T)}.$$

From independence of ΔT from Φ^R and $\tilde{\Phi}$ we know that

$$\mathrm{Cov}(\tilde{\Phi}\,\Delta T,\ \Phi^R\,\Delta T) = \mathrm{Var}(\tilde{\Phi}\,\Delta T) + \mathrm{Cov}(\Phi^R, \tilde{\Phi}) \cdot \mathrm{Var}(\Delta T),$$

which implies

$$\frac{\mathrm{Cov}(\tilde{\Phi}\,\Delta T,\ \Phi^R\,\Delta T)}{\mathrm{Var}(\tilde{\Phi}\,\Delta T)} = 1 + \frac{\mathrm{Cov}(\Phi^R, \tilde{\Phi}) \cdot \mathrm{Var}(\Delta T)}{\mathrm{Var}(\tilde{\Phi}\,\Delta T)} > 1$$

since $\mathrm{Cov}(\Phi^R, \tilde{\Phi}) > 0$ and $\mathrm{Var}(\cdot) > 0$. Hence,

$$\tilde{\beta} = \beta^R \cdot \left(1 + \frac{\mathrm{Cov}(\Phi^R, \tilde{\Phi})\ \mathrm{Var}(\Delta T)}{\mathrm{Var}(\tilde{\Phi}\,\Delta T)}\right) > \beta^R.$$

Thus, since $\tilde{\beta}$, estimated with ϕ^* and $\theta = 1$, exceeds the pass-through coefficient β^R that would be obtained by re-estimating with $\phi = \phi^R$ and $\theta = 1$ (i.e. $\tilde{\beta} > \beta^R$), $\tilde{\beta}$ constitutes an upper bound on the pass-through coefficient in equation (E.4). Consequently, the estimates for $\theta - 1$ of Figure 3—obtained under $\phi = \phi^*$ and $\theta = 1$ via the pass-through formula—can be interpreted as *upper bounds* on the coefficients one would recover by re-estimating with $\phi = \phi^R$ at $\theta = 1$.