Payment Terms and Heterogeneous Shock Transmission in Production Networks

Felipe Leal*
UCLA

March, 2026

Abstract

Supplier cost shocks are not transmitted uniformly across buyers within production networks: pre-shock payment terms predict which relationships absorb the shock through reduced quantities. Using universe-level Chilean firm-to-firm transaction and balance-sheet data (2021–2025), I construct supplier input-cost shocks from net exposure to exchange-rate-driven changes in imported-input prices and estimate heterogeneous pass-through across relationships that differ in their pre-shock trade-credit intensity—a predetermined summary statistic of relationship type. Exploiting within seller–buyer–product match variation over time, I find that a one-standard-deviation shock reduces quantities on cash links by about 3%, while quantities on high-trade-credit links remain essentially unchanged. Prices adjust with a near-zero differential in the pooled specification. Restricting to adverse shocks reveals that the supplier also raises prices for trade-credit buyers while rationing cash buyers on the quantity margin. At the supplier level, profitability deteriorates and leverage rises. A partial-equilibrium model with convex costs and relational contracting provides one interpretation of the quantity patterns and links the estimated elasticities to short-run cost convexity. The results imply that the structure of payment terms within production networks is a first-order determinant of which downstream firms bear the real consequences of upstream cost shocks—a margin that existing network models, which typically abstract from heterogeneity across relationships within a supplier's buyer base, do not capture.

JEL Codes: G32, E32, F41, L14

Keywords: Trade credit; Supplier cost shocks; Exchange rate; Production networks; Relational contracts; Shock transmission; Firm-to-firm trade

*UCLA. Email: fleal@ucla.edu. I am extremely grateful for the constant support and guidance of Saki Bigio, Hugo Hopenhayn, Andy Atkeson and Ariel Burstein. Also, thanks to all the participants in the Trade and Macro seminar at UCLA, FEN UChile and the Central Bank of Chile, for their valuable feedback.

1 Introduction

A growing literature shows that idiosyncratic shocks to individual firms propagate through production networks and contribute to aggregate fluctuations (Acemoglu et al. 2012; Gabaix 2011; Bigio and La'O 2020). Existing production-network models typically abstract from heterogeneity in how a supplier allocates shocks across its downstream buyers. This paper shows that such heterogeneity is first-order. Using universe-level Chilean firm-to-firm transaction data, I document that when a supplier faces an exchange-rate-driven increase in imported-input costs, the transmission of that shock to downstream buyers differs sharply with the pre-existing structure of payment terms: quantities on cash-based links decline substantially, while quantities on trade-credit-intensive links remain essentially unchanged. Pre-shock payment terms—an observable, predetermined feature of each buyer–supplier relationship—predict which downstream firms bear the real consequences of upstream cost shocks and which are insulated.

The empirical strategy exploits universe-level Chilean firm-to-firm transaction and balance-sheet data covering 2021–2025. I construct idiosyncratic supplier-year cost shocks from each firm's pre-period net exposure to imported intermediate goods (net of export revenue) interacted with exchange-rate movements. The preferred specification regresses buyer–supplier–product–year log quantities and log unit prices on the supplier's cost shock and its interaction with the pre-shock trade-credit share of the relationship, controlling for seller–buyer–product and year fixed effects. I interpret pre-shock trade-credit intensity as a predetermined summary statistic of relationship type that predicts how suppliers allocate cost shocks across their buyer base. The relational-contracting literature, where firms optimally protect high-surplus relationships when facing adverse conditions (Macchiavello and Morjaria 2015; Baker et al. 2002), provides one natural interpretation of why TC intensity predicts differential responses, but the empirical contribution does not require a specific micro-foundation: the identifying assumption is that, conditional on the fixed effects, the interaction of the predetermined TC share with the contemporaneous cost shock isolates differential responses across relationship types.

The results are sharp. A one-standard-deviation cost shock reduces quantities on cash links by about 3%, whereas quantities on high-trade-credit links barely move. In the pooled specification, the price interaction with TC intensity is near zero, but restricting to adverse cost shocks reveals a richer two-margin pattern: the supplier raises prices for TC buyers while rationing cash buyers on the quantity margin, consistent with different

competitive and relational constraints operating across buyer types. At the supplier level, the shocks significantly lower return on assets and raise leverage, confirming that the cost increase is economically meaningful and that suppliers absorb part of it in their balance sheets.

These findings establish that pre-shock payment terms predict which buyer–supplier links absorb supplier cost shocks through reduced quantities, revealing that pass-through within production networks depends on relationship type, not only on supplier-level exposure. The quantity margin is where the strongest and most robust heterogeneity appears, and it is the margin that matters most for real activity and downstream propagation. Building on relational-contract theory, I offer one interpretation: trade-credit intensity is a predetermined proxy for relationship type—consistent with evidence that trade credit is more common in longer and higher-markup relationships—and suppliers protect high-value relationships by maintaining their quantities while concentrating quantity reductions on marginal, cash-intensive buyers. I formalize this interpretation with a partial-equilibrium model in which the supplier with convex costs protects high-surplus relationships on the quantity margin. The model maps the estimated quantity elasticities into the short-run convexity of supplier costs and delivers a counterfactual quantifying the cost of removing priority-based allocation.

The rest of the paper is organized as follows. Section 2 situates the paper in the literature on production networks, trade credit, relational contracting, and exchange-rate pass-through. Section 3 describes the Chilean setting, the VAT and balance-sheet data, and key descriptive facts about trade credit in the production network. Section 4 presents the empirical strategy, defines the cost shock, and documents average and heterogeneous effects at the link level, along with robustness and identification checks. Section 5 develops the priority-based reallocation interpretation. Section 6 presents a model that maps into the empirical findings. Section 7 uses the model to quantify the implied cost convexity and the cost of removing priority-based allocation. Section 8 discusses extensions toward a richer quantitative framework. Section 9 concludes.

2 Related Literature

Trade credit is a key layer of financial linkages in production networks. A large literature shows that these financial links transform idiosyncratic shocks into aggregate fluctuations and can either dampen or amplify financial disturbances, depending on the tightness of financial frictions and firms' positions in the network (Acemoglu et al. 2012; Gabaix 2011;

Bigio and La'O 2020; Altinoglu 2021; Bocola and Bornstein 2025). Yet this work remains largely silent on how the relational structure of payment terms shapes the *price–quantity decomposition* of shock pass-through across different buyer–supplier relationships, and on the micro mechanisms behind the reallocation patterns that emerge in transaction-level data.

Production-network models with financial frictions emphasize that interfirm payables and receivables propagate shocks across firms and sectors. Acemoglu et al. (2012) and Gabaix (2011) show that idiosyncratic disturbances at central suppliers can generate sizable aggregate volatility when network linkages are sufficiently asymmetric. Bigio and La'O (2020) study how financial distortions in production networks reduce aggregate efficiency and alter the transmission of shocks across firms. Incorporating trade credit as an additional financial layer further raises the contribution of shocks to aggregate fluctuations: in a calibrated model for the U.S., Altinoglu (2021) finds that allowing for credit linkages substantially increases the share of aggregate volatility explained by financial shocks. In related work, Bocola and Bornstein (2025) show that trade credit relaxes borrowing constraints in normal times but amplifies output drops in crises, amplifying the output costs of the Great Recession. At the micro level, Jacobson and von-Schedvin (2015) document that trade-credit defaults transmit distress upstream and substantially raise suppliers' bankruptcy risk, while Costello (2020) shows that credit-supply shocks at banks propagate downstream when constrained firms cut trade credit to their customers.[1]

A second strand of work focuses on trade credit as relationship finance and relational contracting. Classic evidence shows that trade credit substitutes for bank credit when financial institutions are expensive or constrained, particularly for small and bank-dependent firms (Petersen and Rajan 1997; Maksimovic and Demirgüç-Kunt 2001). More recent contributions emphasize the relational nature of trade credit: it is concentrated in strong, repeated relationships with high markups and high surplus, and it co-moves with measures of relationship importance, suggesting that trade credit is not a simple transactional loan but part of a relational contract (Garcia-Marin et al. 2023; Benguria et al. 2023). A broader literature on relational contracting in supply chains establishes that firms prioritize high-value relationships when allocating scarce resources. Macchiavello and Morjaria (2015) document that Kenyan rose exporters, when hit by a negative supply shock, allocate scarce output differentially across buyers depending on relationship characteristics, with deliveries shaped by the value of the ongoing relationship. Baker et al. (2002) and Levin (2003) provide the

1. See also Chodorow-Reich (2014) on employment effects of bank shocks transmitted through firm financing.

theoretical foundations for such behavior: in repeated interactions sustained by self-enforcing contracts, agents optimally protect relationships that generate high continuation value. In this paper, I use the term *relational priority* to refer to the supplier's willingness to protect a given buyer when facing adverse conditions, and I interpret pre-shock trade-credit intensity as an empirical proxy for this object—motivated by the evidence in Garcia-Marin et al. (2023) and Benguria et al. (2023) that trade credit marks high-surplus, high-trust relationships. Closely related, Cazzaniga et al. (2025) use Brazilian firm-to-firm data to show that trade credit reduces output losses from financial frictions faced by downstream firms, highlighting its stabilizing role for customers. Other recent work documents that trade-credit flows can stabilize supply chains when firms face operating or exchange-rate shocks, especially when suppliers are large and central in the network (Ersahin et al. 2025, 2024; Hardy et al. 2024). These papers point to a potential insurance role for relational trade-credit links but do not directly quantify how it manifests in prices versus quantities along individual buyer–supplier links.

A third line of research studies how trade credit mediates shock transmission, distinguishing between different shock types and positions in the network. Bocola and Bornstein (2025) and Altinoglu (2021) show that trade credit can amplify shocks when financially constrained suppliers reduce trade credit and pass distress on to customers, but can also dampen shocks when financially unconstrained suppliers expand trade credit to buffer downstream firms. Cazzaniga et al. (2025) find that trade-credit-providing suppliers substantially reduce the output losses from financial frictions faced by their customers, emphasizing a stabilizing role in the face of downstream or demand-side shocks. At the same time, the macro literature stresses that shocks hitting trade-credit providers themselves can be particularly harmful, because they simultaneously reduce bank credit and the ability of suppliers to extend trade credit, thereby amplifying downturns (Bocola and Bornstein 2025; Altinoglu 2021; Costello 2020). None of these papers, however, study how such shocks are allocated across relationships of the same supplier that differ in their relational intensity, nor how the adjustment splits into prices versus quantities.

Finally, a large literature on exchange-rate pass-through (ERPT) studies how firms adjust prices in response to exchange-rate-driven cost changes and finds that pass-through is incomplete and varies with market power and invoicing currency (Berman et al. 2012; Gopinath et al. 2010; Amiti et al. 2014). The present paper uses a similar source of cost variation—exchange-rate movements interacted with import exposure—but addresses a different question. The ERPT literature asks how much of a cost shock is passed into

prices at the firm level. This paper asks how the *quantity* adjustment is allocated across buyers within a supplier, conditional on the price response. The finding of incomplete price pass-through is consistent with the ERPT literature; the novel contribution is to show that the real consequences of the shock are distributed across the supplier's buyer base according to the structure of payment terms, through a quantity margin that has not been studied in that literature, and that the supplier uses different adjustment margins for different buyer types depending on the competitive and relational constraints it faces.

This paper contributes to all of these strands by focusing on idiosyncratic upstream input-cost shocks and the price–quantity decomposition of their pass-through across relationships that differ in trade-credit intensity. Using universe-level Chilean firm-to-firm transaction and balance-sheet data (2021–2025), I construct supplier-level cost shocks from net exposure to exchange-rate-driven changes in imported-input prices and estimate heterogeneous pass-through exploiting within seller–buyer–product match variation over time. The main finding is that when a supplier is hit by a cost shock, the strongest heterogeneity in pass-through is in quantities: quantities in cash relationships decline substantially while quantities in trade-credit relationships remain largely unchanged. At the supplier level, these shocks reduce return on assets and increase leverage. I interpret pre-shock trade-credit intensity as an empirical proxy for relational priority, supported by evidence that the TC interaction survives a saturated specification controlling for buyer and supplier size, age, export status, and relationship length, and that the theoretical model provides a monotone mapping from relational surplus to trade-credit intensity. Together, the results show that the structure of payment terms—as captured by trade-credit intensity—is a first-order determinant of how upstream cost shocks are allocated across a supplier's buyer base.

On the theory side, existing dynamic models endogenize trade credit through limited commitment, working-capital constraints, and relational rents, and they deliver both amplification and dampening of shocks at the macro level depending on financial tightness (Cuñat 2007; Bocola and Bornstein 2025; Cazzaniga et al. 2025). However, they do not directly speak to the cross-sectional heterogeneity documented here: they typically abstract from explicit differences across relationships within a supplier and do not reproduce a pattern in which cost shocks generate highly heterogeneous quantity adjustments across buyer types. This paper fills that gap by documenting the empirical patterns and developing a relational model in which the supplier protects high-surplus relationships on the quantity margin while concentrating real adjustment on marginal buyers. The model provides one micro-foundation for what I term *priority-based allocation*—the pattern whereby high-TC relationships are

insulated on the quantity margin from frequent idiosyncratic upstream shocks.

3 Background and Data

3.1 Facts

Trade credit is the dominant source of short-run external financing for firms. Table 1 reports the share of debt over sales by source of financing using the 2019 Longitudinal Survey of Firms (ELE). Across all size classes, supplier financing (trade credit) exceeds bank financing in the short run: for micro firms, short-run debt from suppliers is around ten times larger than short-run bank debt, and even for large firms, supplier credit accounts for more than twice as much short-run debt as bank loans.

Table 1: Share Debt/Sales by Source of Financing

	Short–Run		Long–Run		Total	
Size	Bank	Suppliers	Bank	Suppliers	Bank	Suppliers
Micro	2.6	26.4	3.9	2.7	6.6	29.1
Small	4.3	27.5	14.8	9.2	19.2	36.7
Medium	20.2	51.9	22.2	12.6	42.5	64.5
Large	22.8	56.8	19.1	10.2	41.9	67.1

Source: ELE 2019. Firm size thresholds defined for micro: $<$ 100k (73%); Small: $<$ 1M (21%); Medium: $<$ 4M (4%); Large: $>$ 4M (2%), all in USD.

Turning to the VAT data for the universe of Chilean firms, Table 2 reports descriptive statistics on firm-to-firm trade and trade credit by firm size. On the supply side, large suppliers are only about one-third of all suppliers, but the average large supplier has annual sales of approximately 261 million USD and sells nearly 5 billion USD to other firms, of which about 4.2 billion are on trade-credit terms. On the demand side, a relatively small group of large buyers accounts for the bulk of firm-to-firm purchases and trade-credit use.

Table 3 expresses these magnitudes as shares of totals. Large suppliers account for roughly 98–99% of total firm-to-firm trade and trade-credit volume despite representing only around 32% of suppliers. Large buyers are about 5% of all buyers but generate about 98% of total sales and around 70–80% of firm-to-firm and trade-credit purchases. The Chilean production network is thus organized around a small core of large firms that intermediate the vast majority of interfirm trade and trade credit.

Table 2: Descriptive statistics by size

	Micro	Small	Medium	Large
Suppliers				
Number of firms	33	265	265	272
Sales (M USD)	0.06	0.66	2.23	261
Firm-to-firm trade (M USD)	0.30	16.31	83.25	4,942
TC (M USD)	0.23	12.16	67.09	4,242
Buyers				
Number of firms	101,474	73,584	14,057	7,900
Sales (M USD)	0.04	0.34	1.89	237
Firm-to-firm trade (M USD)	213	821	467	3,536
TC (M USD)	65	483	368	3,403

Note: Sales (M USD) is the firm average within that size group.

Table 3: Descriptive statistics by size (horizontal %)

	micro	small	medium	large
Suppliers				
Number of firms	4	32	32	32
Sales	0.02	0.2	0.8	98.8
Firm-to-firm trade	0.6	0.3	1.7	97.9
TC	0.5	0.3	1.6	98.1
Buyers				
Number of firms	51	37	7	5
Sales	0.1	0.1	0.8	98.0
Firm-to-firm trade	4.2	16.3	9.3	70.1
TC	1.5	11.2	8.5	78.6

Note: percentages calculated over the total across size for each variable.

However, trade-credit intensity is high across almost all size groups, not just for large firms. Table 4 shows that, on the supply side, the share of firm-to-firm sales conducted on credit is already around 75–80% for micro, small, and medium suppliers and rises to about 86% for large suppliers. On the demand side, the gradient is steeper: only micro buyers are predominantly cash-based (31%), whereas small, medium, and large buyers purchase

roughly 59%, 79%, and 96% of their inputs on trade credit, respectively. The combination of uniformly high TC intensity on the supply side and heterogeneous TC intensity on the buyer side means that the same large supplier typically serves both high-TC and low-TC buyers—a feature that is central to the empirical strategy.

Table 4: Share of TC over total firm-to-firm trade (%)

	micro	small	medium	large
Suppliers	76	75	81	86
Buyers	31	59	79	96

Table 5 documents the structure of relationships at the firm level. The median supplier sells to around 104 distinct buyers and 92 products per year, with a trade-credit share tightly concentrated around 80% (standard deviation of 0.04). Buyers, in contrast, source from a median of only 2 suppliers and exhibit a much more dispersed trade-credit share (mean 0.47, standard deviation 0.18). This asymmetry—many buyers per supplier, few suppliers per buyer—implies that buyers are not well diversified across suppliers, so that a cost shock to a single supplier may have a meaningful impact on the affected buyer. Invoice-level payment delays are short on average: the unconditional median is zero (reflecting the large mass of cash transactions), while among invoices with strictly positive delay the median is approximately 30 days.

A natural concern is whether trade credit in this setting primarily reflects mechanical monthly billing cycles rather than genuine deferred payment. To assess this, I examine the extent to which trade-credit invoices between a given buyer–supplier pair cluster on the same day within a month. Table 6 reports the distribution of this clustering share across relationships, measured both by number of invoices and by transaction value. The median share of trade-credit invoices that arrive on the same day within a month is 15% by count and 11% by value, indicating that the majority of trade-credit transactions within a relationship are spread across different days rather than concentrated in a single billing event. At the same time, the distribution is right-skewed: above the 75th percentile, clustering is substantial (around 72–75%), suggesting that monthly billing does account for a non-negligible share of trade credit in some relationships. However, for the network as a whole, same-day within-month clustering accounts for a minority of both the number and value of trade-credit invoices, supporting the interpretation that trade credit in this setting predominantly reflects genuine deferred payment.

Table 5: Descriptive statistics

	mean	sd	p10	p25	p50	p75	p90	obs
Delay	7	15	0	0	0	2	30	898,251
Delay (if > 0)	24	19	1	7	30	30	45	262,165
Suppliers								2,483
N buyers	865	7,650	3	21	104	401	1,157	
N products	1,185	20,021	2	12	92	397	1,209	
TC/total trade	0.80	0.04	0.72	0.76	0.80	0.84	0.86	
Sales (M USD)	27	289	0.24	0.61	1.9	5.8	17	
Buyers								583,013
N suppliers	2	2	1	1	2	3	4	
N products	101	220	5	15	43	106	222	
TC/total trade	0.47	0.18	0.29	0.30	0.31	0.59	0.74	
Sales (M USD)	2.5	79	0.009	0.028	0.09	0.31	1.1	

Table 6: Clustering of trade-credit invoices within month

	Mean	SD	p5	p25	p50	p75	p95
Share of invoices	0.34	0.38	0.00	0.00	0.15	0.75	0.98
Share of value	0.33	0.38	0.00	0.00	0.11	0.72	0.98

Taken together, the descriptive evidence establishes several features of the Chilean production network that are central to the empirical strategy. First, trade credit is the main source of short-run firm financing, exceeding bank credit for all size classes (Table 1), and it is the predominant mode of firm-to-firm trade for all but the smallest buyers (Table 4). Second, both the volume and the provision of trade credit are highly concentrated: roughly 32% of suppliers account for over 98% of trade-credit volume (Table 3), and the average large supplier transacts billions of USD with other firms (Table 2). Third, while suppliers' trade-credit shares are tightly distributed around 80%, buyer-level trade-credit shares are much more dispersed (mean 0.47, standard deviation 0.18, Table 5), implying that the same large supplier typically has both high-TC and low-TC buyers. Fourth, buyers are only moderately diversified across suppliers—the median buyer sources from just two suppliers (Table 5)—so that a cost shock to a single supplier is not easily diversified away. Fifth, trade credit predominantly reflects genuine deferred payment rather than mechanical billing cycles (Table 6). This structure—large suppliers with heterogeneous buyers, limited buyer diversification, and widespread genuine trade credit—makes the Chilean VAT network a suitable setting to study how idiosyncratic supplier cost shocks are transmitted across relationships that differ in their

trade-credit intensity.

4 Empirical Strategy

This section describes the construction of the supplier cost shock and establishes its basic properties. It then examines the impact of the shock on suppliers' financial conditions. Next, it presents the main link-level specification, documenting that quantity adjustments differ substantially across relationships with different trade-credit intensity, and provides robustness checks. It then examines how the adjustment pattern varies across adverse and favorable cost shocks. Finally, it discusses the interpretation of trade-credit intensity as a predetermined marker of relationship type.

4.1 Shock construction and basic properties

For each supplier $j \in \mathcal{J}$ and year t, I construct an idiosyncratic input-cost shock ν_{jt} that exploits variation in each supplier's *net* exposure to exchange-rate-driven changes in the cost of imported intermediates. The logic is that a supplier importing a large share of its inputs is more affected by a peso depreciation than one sourcing domestically, but that exporting suppliers are partially hedged because the depreciation also raises the local-currency value of their export revenue (Amiti et al. 2014).

For every supplier j, define the net import exposure as

$$\text{exposure}_{j,t} = \frac{\text{imports of intermediate inputs}_{j,t} - \text{exports}_{j,t}}{\text{total expenditure on intermediate inputs}_{j,t}}.$$

This measure captures the share of the firm's input costs that is effectively exposed to exchange-rate movements after accounting for the natural hedge provided by export revenue. The shock is then constructed as

$$\nu_{jt} = \text{exposure}_{j,t-1} \times \Delta \log \text{ER}_t,$$

where $\Delta \log \text{ER}_t$ is the annual log change in the Chilean peso–dollar exchange rate. The exposure is lagged one period to avoid contemporaneous endogeneity: the supplier's import intensity as of the year before the exchange-rate movement is used, so that the shock reflects pre-determined compositional exposure interacted with an aggregate price change.[2]

2. The shock used in estimation is the residual from regressing ν_{jt} on supplier controls and sector–location–

Pre-shock trade-credit intensity. To characterize the pre-existing relational structure of each buyer–supplier pair, I define a supplier–buyer–year measure of pre-shock trade-credit intensity between supplier j and buyer i:

$$\mathrm{TC}_{ij,t}^{\mathrm{pre}} = \frac{1}{2}\sum_{l=1}^{2} \frac{\sum_b \mathrm{TC}_{ijb,t-l}}{\sum_b \mathrm{TC}_{ijb,t-l} + \mathrm{Cash}_{ijb,t-l}},$$

where $\mathrm{TC}_{ijb,t-l}$ denotes the value of transactions between supplier j and buyer i in product b conducted on trade-credit terms in year $t-l$, and $\mathrm{Cash}_{ijb,t-l}$ denotes the corresponding cash transactions. The measure aggregates across all products b traded between the pair and averages over the two preceding years ($l = 1, 2$). Averaging over two lags serves two purposes: it smooths year-to-year fluctuations in the payment-term mix, and it captures the established relational character of the link rather than transitory changes in contractual terms. A relationship with $\mathrm{TC}_{ij,t}^{\mathrm{pre}}$ close to one is one where the supplier has consistently extended credit to the buyer over the recent past; a relationship near zero is conducted primarily on a cash basis.

Orthogonality of trade-credit intensity and the shock. A requirement for interpreting heterogeneous responses by $\mathrm{TC}_{ij,t}^{\mathrm{pre}}$ is that pre-shock trade-credit intensity is not systematically correlated with the shock itself. If, for instance, suppliers with more trade-credit-intensive buyer bases also tended to have lower import exposure, the interaction could confound relational heterogeneity with differential shock magnitude. To assess this, I examine the relationship between $\mathrm{TC}_{ij,t}^{\mathrm{pre}}$ and both the level and absolute value of ν_{jt} at the supplier–buyer–year level. The pre-shock trade-credit share explains around 0.25% of the variation in $|\nu_{jt}|$ and approximately 1% of the variation in ν_{jt}. These negligible R^2 values indicate that more trade-credit-intensive relationships are neither systematically associated with suppliers facing larger shocks nor with suppliers receiving more favorable shocks. The interaction term in the main specification therefore identifies differential *responses* to a given shock rather than differential *exposure* to shocks of different sizes or signs.

---

year fixed effects:

$$\nu_{jt} = \beta \mathsf{X}_j + \gamma_{s\times l\times t} + \varepsilon_{j,t},$$

where $\gamma_{s\times l\times t}$ is a 6-digit sector–location–year fixed effect and X_j includes the lagged shock, firm age, and export status. This residualization ensures that the identifying variation is idiosyncratic to the supplier after absorbing sector–location–year trends and observable firm characteristics.

4.1.1 Supplier-year financial outcomes

Before turning to the link-level analysis, I verify that the constructed cost shock has economically meaningful effects on suppliers' balance sheets. I estimate:

$$f_{j,t} = \theta \nu_{j,t} + \alpha_j + \gamma_t + \mu_{j,t},$$

where $f_{j,t}$ is a financial outcome, α_j are supplier fixed effects, and γ_t are year fixed effects; standard errors are clustered at the supplier level. I consider three outcomes. Return on assets, $\text{ROA}_{j,t} = \text{Net Income}_{j,t}/\text{Assets}_{j,t}$, measures how efficiently the firm converts its asset base into profits. Leverage, $\text{Leverage}_{j,t} = \text{Liabilities}_{j,t}/\text{Assets}_{j,t}$, measures the firm's reliance on external financing relative to its asset base. The profit margin, $\text{Profit Margin}_{j,t} = \text{Net Income}_{j,t}/\text{Sales}_{j,t}$, measures the share of revenue that translates into net income.

Table 7 reports the results for large firms, who as suppliers account for the vast majority of firm-to-firm sales and trade-credit volume (Table 3). A higher cost shock ν_{jt} reduces ROA by 0.111 and increases leverage by 0.099, both statistically significant. These estimates indicate that the shock deteriorates suppliers' balance sheets: returns on assets decline as profitability falls, while liabilities relative to assets rise as suppliers increase their reliance on external financing to cover increased input costs. The coefficient on the profit margin is small, negative (-0.025), and statistically insignificant, suggesting no large change in the net-income-to-sales ratio. This pattern is consistent with the cost increase being absorbed partly in lower earnings and partly in higher leverage, without a large contraction in operating margins.[3]

3. For non-large suppliers, ROA and leverage responses are smaller in magnitude but also statistically significant (-0.068 and 0.062, respectively), while the profit margin increases with the shock (2.380, Appendix Table 14). The smaller magnitudes are consistent with non-large suppliers having lower import exposure on average. The core conclusion—that the cost shock has real balance-sheet consequences—holds across the firm size distribution.

Table 7: Supplier-year financial outcomes: large suppliers

	dep. variable:		
	ROA_{jt}	Leverage_{jt}	$\text{Profit Margin}_{jt}$
ν_{jt}	-0.111***	0.099***	-0.025
	(0.017)	(0.015)	(0.073)
supplier FE	Yes	Yes	Yes
year FE	Yes	Yes	Yes
obs	125,146	126,657	126,657

Note: Large suppliers as defined in Table 2. Standard errors clustered at the supplier level.

4.2 Main link-level specification

The main analysis examines buyer–supplier–product–year outcomes. Denote by y_{ijbt} a link-level outcome—either the log quantity purchased by buyer i from supplier j of product b in year t, or the corresponding log unit price (computed as total transaction value divided by total quantity across invoices within the match-year). Let $\text{TC}_{ij,t}^{\text{pre}}$ denote the pre-shock trade-credit share defined above. The preferred specification is:

$$y_{ijbt} = \beta_1 \nu_{jt} + \beta_2 \text{TC}_{ij,t}^{\text{pre}} + \beta_3 \nu_{jt} \times \text{TC}_{ij,t}^{\text{pre}} + \alpha_{ijb} + \gamma_t + \varepsilon_{ijbt}, \tag{1}$$

where α_{ijb} are seller–buyer–product fixed effects and γ_t are year fixed effects; standard errors are clustered at the supplier level. The seller–buyer–product fixed effects absorb all time-invariant heterogeneity at the level of a specific buyer–supplier–product match, including relationship quality, geographic proximity, and product-specific characteristics. Year fixed effects absorb aggregate shocks common to all matches in a given year. Identification therefore exploits *within-match variation over time*: when supplier j is hit by a cost shock in year t, the specification compares how the quantity or price in each of j's matches changes relative to that match's own time-invariant level. The coefficient β_3 is identified because different matches of the same supplier have different pre-shock TC shares, so that the same supplier-level shock ν_{jt} generates differential responses across matches according to their relational status.

In this specification, β_1 captures the effect of the shock on matches with zero pre-shock trade credit ("cash" relationships), β_2 captures the within-match time-varying association between TC intensity and the outcome, and $\beta_1 + \beta_3 \times \text{TC}_{ij,t}^{\text{pre}}$ captures the effect of the shock

on matches at a given level of pre-shock trade-credit intensity. Because y_{ijbt} is in logs and ν_{jt} is proportional to the percentage change in imported-input costs, β_1 and β_3 can be interpreted as elasticities with respect to that cost component.

Table 8: Main link-level specification: quantities and prices

	dep. variable:	
	$\log(\text{quantity}_{ijbt})$	$\log(\text{price}_{ijbt})$
ν_{jt}	-1.492***	0.440**
	(0.446)	(0.203)
$\nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t}$	1.475***	-0.093
	(0.302)	(0.184)
seller–buyer–product FE	Yes	Yes
year FE	Yes	Yes
obs	34,518,969	34,518,969

Note: standard errors clustered at the supplier level.

Table 8 reports the results. For quantities, $\hat{\beta}_1 = -1.492$ and $\hat{\beta}_3 = 1.475$. A 1% increase in the imported-input cost component is associated with a 1.5% decline in quantities for pure cash links, while the positive interaction term progressively offsets this decline for higher-TC links. At the median pre-shock TC share of approximately 0.80, the net quantity effect is $\hat{\beta}_1 + \hat{\beta}_3 \times 0.80 \approx -0.31$, a modest decline. At $\text{TC}^{\text{pre}} = 1$, the net effect is $\hat{\beta}_1 + \hat{\beta}_3 \approx -0.02$, essentially zero. Using the observed standard deviation of ν_{jt}, a one-standard-deviation shock is associated with a decline in cash-link quantities of about 3%, while quantities on high-TC links barely change. The gap in quantity responses between cash and high-TC relationships is approximately 3 percentage points per one-standard-deviation shock.

For prices, the level coefficient ($\hat{\beta}_1 = 0.440$) is positive and statistically significant, indicating that suppliers do pass through part of the cost increase into unit prices. The interaction ($\hat{\beta}_3 = -0.093$) is small and statistically insignificant, indicating that, on average across all shock realizations, the price response does not differ significantly across relationship types. At the median TC share, the price effect is $0.440 + (-0.093) \times 0.80 \approx 0.37$, similar to the effect on cash links. Section 4.4 examines whether this average masks state-dependent pricing patterns.

The decomposition in the pooled specification is therefore clear: the price interaction is near zero while the quantity interaction is large and significant. Cash links face a substantial

quantity reduction; high-TC links face essentially no quantity reduction. The strongest and most robust heterogeneity in pass-through operates through quantities.

Two features of this decomposition deserve emphasis. First, the quantity reduction on cash links should not be understood as a movement along the demand curve in response to the price increase. In the pooled specification, the price interaction is near zero, so the price response is approximately common across buyer types. A 0.44% price increase, even with a demand elasticity as high as 3–4, would generate a quantity decline on the order of 1.3–1.8%—and it would apply symmetrically to all buyers, not differentially to cash links. The observed cash-link quantity decline exceeds what the price increase alone could produce, and the 3 percentage-point gap between cash and high-TC links cannot be explained by a price response that is approximately common in the pooled specification. The differential quantity response therefore reflects a supply-side allocation decision, not differential demand responses to a common price change.

Second, the combination of differential quantity responses and balance-sheet deterioration points to a specific mechanism: *supplier-side rationing*. The supplier faces a cost increase and responds on two margins. On the price margin, it raises the price slightly for all buyers—recovering part of the cost increase, but not enough to fully offset it (as confirmed by the decline in ROA and the increase in leverage). On the quantity margin, the supplier's optimal total output is lower than before the shock, because at the old output level marginal cost now exceeds the price. The supplier must therefore allocate a reduced quantity across its buyer base. It fills orders for high-TC buyers at their pre-shock level, absorbing the higher per-unit cost on these deliveries in its margins and balance sheet, because the value of preserving these high-surplus relationships (J^{TC}) exceeds the short-run cost of maintaining their quantities. Cash buyers receive the residual: the supplier reduces their shipments, because the relational penalty for under-serving a cash buyer is negligible. The cash buyer receives less not because it chose to buy less at the higher price, but because the supplier allocated less to it.

4.3 Robustness

4.3.1 Average link-level effects

A natural concern is whether the heterogeneous effects documented above could be driven by large average movements in quantities or prices. To address this, I estimate a specification

identical to the preferred one but without the trade-credit interaction:

$$y_{ijbt} = \beta_1 \nu_{jt} + \alpha_{ijb} + \gamma_t + \varepsilon_{ijbt},$$

with seller–buyer–product and year fixed effects. This specification asks: on average across all of a supplier's matches, does the cost shock move quantities or prices?

Table 9 reports the results. The point estimate on log quantity (−0.846) is negative but statistically insignificant, indicating that the cost shock does not generate a large average contraction in link-level quantities. In contrast, the point estimate on log price (0.440) is positive and statistically significant, confirming that on average across all of a supplier's matches, the cost shock raises unit prices. This average-effects specification reinforces the decomposition from the main results: the average price adjusts significantly while the average quantity response is muted, consistent with within-supplier reallocation across relationship types being the principal quantity margin rather than a uniform contraction.

Table 9: Average effects

	dep. variable:	
	$\log(\text{quantity}_{ijp,t})$	$\log(\text{price}_{ijp,t})$
ν_{jt}	-0.846	0.440**
	(0.529)	(0.171)
supplier-buyer-product FE	Yes	Yes
year FE	Yes	Yes
obs	45,547,979	45,547,979

Note: clusters at supplier level.

4.3.2 Alternative fixed effects

I consider two alternative fixed-effects structures.[4]

The first replaces the seller–buyer–product fixed effect with separate buyer–product and seller–product fixed effects, while adding product–year fixed effects:

$$y_{ijbt} = \beta_1 \nu_{jt} + \beta_2 \text{TC}^{\text{pre}}_{ij,t} + \beta_3 \nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t} + \alpha_{ib} + \delta_{jb} + \gamma_{bt} + \varepsilon_{ijbt}.$$

4. The alternative fixed-effects specifications reported in this subsection use the gross import-exposure measure (imports of intermediate inputs divided by total intermediate inputs). The Appendix documents that the quantity interaction is stable across all exposure measures.

The buyer–product fixed effects α_{ib} absorb all time-invariant buyer–product heterogeneity. The seller–product fixed effects δ_{jb} absorb all time-invariant supplier–product heterogeneity. The product–year fixed effects γ_{bt} absorb product-specific aggregate trends, such as changes in commodity prices or shifts in aggregate demand for a particular product category. Identification in this specification comes from cross-match variation: for a given product in a given year, how do different buyer–supplier pairs exposed to different supplier shocks and different TC shares respond? As reported in Table 10, the interaction term on quantities remains positive and significant (1.464), with a magnitude close to the baseline. The level effect on quantities (−1.488) is also similar. For prices, the level coefficient is positive and marginally significant (0.440), while the interaction remains small and insignificant (−0.092), confirming the pattern of a near-zero average price differential across relationship types. The stability of the estimates under this different fixed-effects structure rules out the possibility that the baseline results are driven by product-specific trends or by the particular structure of the triple fixed effect.

Table 10: Alternative FE: buyer–product, seller–product, product–year

	dep. variable:	
	$\log(\text{quantity}_{ijp,t})$	$\log(\text{price}_{ijp,t})$
ν_{jt}	-1.488***	0.440*
	(0.568)	(0.258)
$\nu_{jt} \times \text{TC}_{ij,t}^{\text{pre}}$	1.464***	-0.092
	(0.386)	(0.234)
buyer-product FE	Yes	Yes
seller-product FE	Yes	Yes
year FE	Yes	Yes
obs	34,800,296	34,800,296

Note: clusters at supplier level.

The second specification restores the seller–buyer–product fixed effect and adds product–year fixed effects:

$$y_{ijbt} = \beta_1 \nu_{jt} + \beta_2 \text{TC}_{ij,t}^{\text{pre}} + \beta_3 \nu_{jt} \times \text{TC}_{ij,t}^{\text{pre}} + \alpha_{ijb} + \gamma_{bt} + \varepsilon_{ijbt}.$$

This is the most demanding specification in the paper. The seller–buyer–product fixed effects absorb all time-invariant match-level heterogeneity, and the product–year fixed effects absorb all product-specific aggregate variation. Identification therefore relies on within-match, within-

product–year variation. As reported in Table 11, the interaction term on quantities remains positive and statistically significant (0.228), though attenuated relative to the baseline. The attenuation is expected: the product–year fixed effects absorb a substantial share of the variation in ν_{jt} that is common across products within a year. Accordingly, the main effect of the shock (-0.147) becomes smaller and statistically insignificant, a mechanical consequence of the product–year fixed effects absorbing the average component of the shock. The key result is that the interaction survives: the differential quantity response across relationship types is robust even under this stringent identification. Price coefficients show a positive but insignificant level (0.274) and a near-zero interaction (-0.029), consistent with the baseline pattern of a near-zero average price differential.

Table 11: Alternative FE: seller–buyer–product, product–year

	dep. variable:	
	$\log(\text{quantity}_{ijp,t})$	$\log(\text{price}_{ijp,t})$
ν_{jt}	-0.147	0.274
	(0.224)	(0.185)
$\nu_{jt} \times \text{TC}_{ij,t}^{\text{pre}}$	0.228**	-0.029
	(0.111)	(0.026)
seller-buyer-product FE	Yes	Yes
product-year FE	Yes	Yes
obs	34,210,462	34,210,462

Note: clusters at supplier level.

4.3.3 Slope controls and observed heterogeneity

A remaining concern is that $\text{TC}_{ij,t}^{\text{pre}}$ might proxy for other observable dimensions of heterogeneity. To address this, I estimate a specification that interacts ν_{jt} with a comprehensive set of supplier and buyer observables—size dummies for both sides of the match, firm age, export status, and the length of the buyer–supplier relationship—in addition to the trade-credit interaction $\nu_{jt} \times \text{TC}_{ij,t}^{\text{pre}}$. This specification includes seller–buyer–product and year fixed effects.

Table 12 reports the results. The trade-credit interaction on quantities remains economically large and statistically significant (0.967), confirming that the differential quantity response is associated with the trade-credit dimension of the relationship even after controlling for a comprehensive set of observable heterogeneity. Relationship length—which could independently proxy for relationship quality—has a small and statistically insignificant interaction (0.065), suggesting that trade-credit intensity captures something specific about

the relationship that duration alone does not.

Several buyer and supplier size interactions are also statistically significant. On the buyer side, the interaction with small-buyer status is negative and significant (-0.402), while the interaction with large-buyer status is positive and significant (0.676), indicating that larger buyers are more insulated from supplier cost shocks, plausibly through greater bargaining power or better outside options. The trade-credit interaction, however, survives at 0.967 after conditioning on these buyer-size effects, indicating that the relational dimension captured by TC intensity operates above and beyond buyer size.

The main effect of the shock on quantity (-10.123) is large, reflecting the effect for the omitted baseline group (micro suppliers at zero values of all interacted variables). The supplier-size interactions (ranging from 8.419 to 8.798, all statistically significant) progressively offset this large baseline. For large suppliers—who account for over 98% of trade-credit volume—the effective main effect is approximately $-10.123 + 8.419 \approx -1.7$, and the TC interaction of 0.967 offsets a substantial share of this effect at the median TC share, consistent with the magnitudes in the baseline specification.

Table 12: Slope controls and observed heterogeneity

	dep. variable:	
	$\log(\text{quantity}_{ijp,t})$	$\log(\text{price}_{ijp,t})$
ν_{jt}	-10.123**	7.382**
	(4.303)	(2.863)
$\nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t}$	0.967***	-0.472
	(0.367)	(0.347)
$\nu_{jt} \times \text{age}_{j,t}$	0.050	-0.012
	(0.050)	(0.022)
$\nu_{jt} \times \text{age}_{i,t}$	-0.034*	-0.007
	(0.019)	(0.007)
small=$1_{j,t}$	0.154	0.090***
	(0.095)	(0.018)
medium=$1_{j,t}$	0.148	0.095***
	(0.104)	(0.025)
large=$1_{j,t}$	0.165	0.122***
	(0.114)	(0.032)
$\nu_{jt} \times$ small=$1_{j,t}$	8.722**	-7.126**
	(4.174)	(2.786)
$\nu_{jt} \times$ medium=$1_{j,t}$	8.798**	-6.295**
	(4.172)	(2.787)
$\nu_{jt} \times$ large=$1_{j,t}$	8.419**	-6.453**
	(4.198)	(2.791)
small=$1_{i,t}$	0.247***	0.000
	(0.039)	(0.001)
medium=$1_{i,t}$	0.405***	0.002
	(0.069)	(0.002)
large=$1_{i,t}$	0.530***	-0.000
	(0.092)	(0.003)
$\nu_{jt} \times$ small=$1_{i,t}$	-0.402***	0.002
	(0.155)	(0.046)
$\nu_{jt} \times$ medium=$1_{i,t}$	-0.055	-0.005
	(0.223)	(0.082)
$\nu_{jt} \times$ large=$1_{i,t}$	0.676**	-0.009
	(0.331)	(0.142)
export=$1_{j,t}$	0.082***	0.044***
	(0.018)	(0.011)
$\nu_{jt} \times$ export=$1_{j,t}$	-0.534	-0.177
	(0.531)	(0.300)
export=$1_{i,t}$	0.004	-0.004
	(0.026)	(0.003)
$\nu_{jt} \times$ export=$1_{i,t}$	0.464	0.319***
	(0.406)	(0.117)
$\nu_{jt} \times$ relationship length (years)$_{ij,t}$	0.065	0.013
	(0.137)	(0.064)
seller-buyer-product FE	Yes	Yes
year FE	Yes	Yes
obs	34,518,969	34,518,969

Note: clusters at supplier level.

4.4 State-dependent adjustment: adverse versus favorable shocks

The pooled specification averages across both adverse (cost-increasing) and favorable (cost-decreasing) shocks. To understand the mechanism more precisely, I restrict the sample to adverse cost shocks ($\nu_{jt} > 0$), where the supplier's capacity is constrained and the rationing motive is strongest.

Table 13: Adverse cost shocks only ($\nu_{jt} > 0$)

	dep. variable:	
	$\log(\text{quantity}_{ijbt})$	$\log(\text{price}_{ijbt})$
ν_{jt}	-1.037***	-0.321
	(0.336)	(0.520)
$\nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t}$	1.347***	0.731**
	(0.380)	(0.297)
seller–buyer–product FE	Yes	Yes
year FE	Yes	Yes
obs	34,519,734	34,519,734

Note: standard errors clustered at the supplier level.

Table 13 reports the results. The quantity estimates confirm the baseline pattern: the main effect on cash links is negative and significant (-1.037), and the TC interaction is positive and significant (1.347). At TC $= 0.80$, the net quantity effect is approximately zero. The quantity insulation of high-TC links is driven by adverse-state observations, consistent with the idea that the rationing motive operates when the supplier's capacity is constrained.

The price column reveals a richer pattern than the pooled specification. Under adverse shocks, the price level for cash links is approximately zero (-0.321)—consistent with the supplier being unable to raise cash-buyer prices when competitors' costs are unchanged and the buyer has outside options. The price interaction for TC links, however, is positive and significant (0.731): the supplier raises prices for TC buyers by approximately $-0.321 + 0.731 \times 0.80 \approx 0.26\%$. This pattern is consistent with TC buyers facing higher switching costs—reflecting the value of credit terms, reliable delivery, and after-sale support—which may allow the supplier to raise prices without losing these buyers.

This adverse-shock pattern is consistent with a framework in which different competitive and relational constraints operate across buyer types. TC buyers, who may face higher

switching costs, may accept modest price increases; cash buyers, who likely have access to competitors at the unchanged market price, are harder to charge more but easier to ration. The supplier appears to adjust on the margin where it faces fewer constraints for each buyer type: prices for TC buyers, quantities for cash buyers.

The fact that TC buyers face higher prices under adverse shocks, if anything, strengthens the evidence for supply-side quantity protection. The demand-side effect of a differential price increase on TC links would push their quantities *down* relative to cash links. Yet TC quantities are flat while cash quantities fall, implying that the supply-side quantity protection is even larger than the pooled estimates suggest.

In the pooled specification, the near-zero average price interaction (-0.093) reflects a cancellation: the supplier charges TC buyers somewhat more under adverse shocks but shares cost decreases preferentially with TC buyers under favorable shocks. The quantity insulation, however, is concentrated in adverse states and does not cancel across shock signs.

4.5 What the results weigh against

The combined evidence from the main specification and the robustness exercises rules out several natural alternative explanations for the documented patterns.

A first alternative is that more trade-credit-intensive suppliers receive systematically different shocks. As documented above, pre-shock trade-credit intensity explains less than 1% of the variation in the shock, ruling out the possibility that the interaction captures differential shock exposure rather than differential shock responses.

A second alternative is that the cost shock generates large average quantity movements, and that the interaction merely captures a nonlinearity in this average response. The average-effects specification shows that the shock has no statistically significant effect on average link-level quantities (Table 9), confirming that the dominant quantity margin is within-supplier reallocation across relationship types rather than a uniform contraction.

A third concern is that the constructed shock might represent measurement noise rather than an economically meaningful disturbance. The supplier-level financial regressions rule this out: the shock significantly reduces return on assets and increases leverage for large suppliers (Table 7), demonstrating that it has real balance-sheet consequences.

A fourth alternative is that the differential quantity response is driven entirely by demand-side responses to differential pricing. In the pooled specification, the price interaction is near zero (-0.093), ruling out average differential pricing. Under adverse shocks, TC buyers do

face somewhat higher prices (0.731 interaction), which could in principle reduce TC quantities through a demand channel. However, this alternative predicts the opposite of what the data show: higher TC prices should push TC quantities *down*, not keep them flat. The fact that TC quantities are unchanged despite facing higher prices under adverse shocks implies that the supply-side quantity protection is even stronger than the pooled estimates suggest.

A fifth concern is that the interaction reflects static, non-relational heterogeneity—for example, that cash buyers happen to be systematically smaller, younger, or more sensitive to supplier shocks for reasons unrelated to the relational structure of the relationship. The slope-controls specification addresses this directly: the trade-credit interaction remains significant (0.967) after controlling for interactions with buyer and supplier size, age, export status, and relationship length (Table 12). Some buyer-size interactions are also statistically significant—larger buyers are more insulated, smaller buyers less so—but the trade-credit interaction survives above and beyond these size effects, indicating that the relational dimension captured by TC intensity is not reducible to buyer size. Moreover, the seller–buyer–product fixed effects in the baseline specification absorb all time-invariant match-level characteristics.

4.6 Interpreting the role of trade-credit intensity

The pre-shock trade-credit share $\mathrm{TC}_{ij,t}^{\mathrm{pre}}$ is an equilibrium outcome of the buyer–supplier relationship. It reflects the joint decisions of both parties and is shaped by factors such as trust, match quality, the supplier's cost advantage in providing short-term finance, and the buyer's outside financing options. The identifying assumption for β_3 is therefore *not* that trade-credit intensity is randomly assigned. Rather, the assumption is that, conditional on the seller–buyer–product fixed effect and year fixed effect, the interaction of the predetermined TC share with the contemporaneous cost shock isolates differential quantity responses across relationship types. The estimates should be interpreted as showing that the dimension of match heterogeneity captured by pre-shock payment terms predicts how suppliers allocate cost shocks across their buyer base, rather than as the causal effect of an exogenous increase in trade credit.

Three features of the data support the interpretation of $\mathrm{TC}_{ij,t}^{\mathrm{pre}}$ as a plausible empirical proxy for the priority that a supplier assigns to a given buyer. First, trade-credit intensity is predetermined: it is measured as the average TC share over years $t-1$ and $t-2$, before the cost shock at t is realized. For the TC share to bias β_3, the buyer and supplier would have needed to adjust payment terms one to two years in advance in anticipation of a future exchange-rate movement interacted with the supplier's specific import exposure. This is implausible.

Second, the trade-credit interaction survives a saturated slope-controls specification that includes interactions of the shock with buyer and supplier size, age, export status, and relationship length (Table 12). While some buyer-size interactions are also statistically significant—reflecting that larger buyers are more insulated from supplier shocks, plausibly through greater bargaining power—the trade-credit interaction remains significant after conditioning on these size effects. This indicates that TC captures a relational dimension that is distinct from buyer size. Relationship length, which could independently proxy for relationship quality, has a small and statistically insignificant interaction (0.065), suggesting that trade-credit intensity captures something specific about the relationship that duration alone does not. Third, the theoretical model developed below provides a monotone mapping from relational surplus to trade-credit intensity, consistent with the empirical findings of Garcia-Marin et al. (2023) and Benguria et al. (2023) that trade credit is concentrated in high-markup, high-surplus relationships.

5 Interpretation: Priority-Based Reallocation

The empirical results establish a robust fact: pre-shock trade-credit intensity predicts how suppliers allocate the consequences of cost shocks across their buyer base. The quantity margin is where the strongest and most robust heterogeneity appears, and it is the margin that matters most for real activity and downstream propagation. This section discusses possible interpretations.

The cost shock significantly reduces return on assets and increases leverage for large suppliers (Table 7), indicating that input-cost increases tighten suppliers' balance sheets. At the link level, the supplier concentrates the quantity adjustment on cash relationships: quantities on cash links decline substantially ($\hat{\beta}_1 = -1.492$), while the interaction with TC intensity is large, positive, and significant ($\hat{\beta}_3 = 1.475$), so that quantities on high-TC links are largely insulated (Table 8). The average quantity effect is muted (-0.846), consistent with within-supplier reallocation rather than uniform contraction (Table 9). Under adverse shocks, the supplier also raises prices for TC buyers while leaving cash-buyer prices unchanged (Table 13), suggesting that the supplier uses different adjustment margins for different buyer types depending on the competitive and relational constraints it faces.

A natural interpretation, grounded in relational-contract theory, is that the supplier responds to the cost shock on the margin where it faces fewer constraints for each buyer type. TC buyers, who may face higher switching costs due to the value of relational services, may accept modest price increases rather than lose the relationship; but the supplier has an

incentive to maintain their quantities, because short-shipping would damage the relationship and threaten future surplus. Cash buyers, who likely have access to competitors at the unchanged market price, are harder to charge more; but the supplier may ration their shipments more freely, because there is less relational surplus at stake. Pre-shock payment terms enter as a state variable that predicts the type of adjustment each buyer experiences.

This interpretation is consistent with models in which trade credit arises endogenously as part of a relational contracting problem. Garcia-Marin et al. (2023) show that trade credit is more likely to be extended in relationships with higher markups and where suppliers have a cost advantage in providing short-term finance relative to banks, and that trade credit use increases with relationship length as parties learn about each other. Benguria et al. (2023) document that trade credit is concentrated in strong, repeated relationships and co-moves with measures of relationship importance. In this framework, relationships with high pre-shock trade-credit intensity are precisely those in which the supplier has accumulated relational capital and has an incentive to protect the relationship when a cost shock arrives.

Relational-contract models explain why cost-shocked suppliers protect some relationships more than others. In Macchiavello and Morjaria (2015), a negative supply shock to Kenyan rose exporters leads suppliers to reallocate quantities across buyers, with high-surplus relationships experiencing smaller quantity reductions than low-surplus relationships, even though the shock is purely on the supplier side. Malcomson (2015) formalizes how relational incentive contracts adjust to productivity shocks while preserving long-term surplus, with the supplier choosing state-contingent allocation rules that depend on relationship value. Translating those insights to this setting, when a supplier's costs rise, it faces a trade-off: it can maintain quantities for some buyers and absorb more of the cost increase in lower markups and higher working-capital needs, or it can reduce quantities more aggressively for other buyers to restore profitability and reduce financing needs. Relationships with high pre-shock trade-credit intensity indicate high relational value and a history of supplier-provided finance, so the optimal policy is to protect those links in adverse states and let quantity adjustments fall disproportionately on cash-intensive relationships.

The empirical patterns are consistent with this interpretation. The firm-level regressions show that cost shocks reduce return on assets and increase leverage, indicating that suppliers absorb part of the cost increase in their margins and balance sheets. Within a supplier's matches, quantities decline much more in relationships with low pre-shock trade-credit shares

than in relationships with high pre-shock trade-credit shares. Under adverse shocks, the supplier raises prices for TC buyers but not for cash buyers, consistent with the supplier facing different competitive constraints across buyer types. The quantity margin is where the strongest and most robust heterogeneity appears, and it is the margin that matters most for real activity and downstream propagation: it determines which firms receive the inputs they need and which do not.

It is worth emphasizing that this interpretation—relational contracting with priority-based allocation—is one plausible micro-foundation for the documented patterns. Alternative mechanisms, such as generic buyer importance, bargaining power, or product criticality, could also generate differential quantity responses correlated with TC intensity. The slope-controls specification shows that the TC interaction survives after conditioning on buyer and supplier size, age, export status, and relationship length, which addresses the observable dimensions of this concern. The relational-contracting framework is offered as the most parsimonious interpretation that is consistent with the full set of empirical patterns—including the state-dependent pricing and the asymmetric quantity response concentrated in adverse states—rather than as the only possible explanation.

Finally, this interpretation is consistent with recent work on trade credit and supply-chain stability. Ersahin et al. (2024) show that trade-credit flows increase when firms in a production network face operating shocks and that this additional supplier financing helps preserve sales and relationships. In this paper's setting, changes in payment terms around shocks are not directly observed, but the cross-sectional evidence shows that relationships with higher pre-shock trade-credit shares are those whose quantities are stabilized, while low-trade-credit (cash) links bear most of the quantity adjustment. Pre-shock trade-credit intensity thus acts as a state variable summarizing relational priority and the supplier's willingness to use its own balance sheet to buffer the shock for particular buyers.

6 A Simple Model of Priority-Based Allocation

This section develops a partial-equilibrium model that provides one micro-foundation for the documented patterns: when facing cost increases, a supplier with convex costs and valuable relational contracts protects high-surplus relationships on the quantity margin while concentrating real adjustment on marginal buyers.

The model is designed to match the central quantity patterns in the data. The first is

supplier balance-sheet deterioration: a positive cost shock reduces ROA ($\hat{\theta}^{ROA} = -0.111$) and increases leverage ($\hat{\theta}^{Lev} = 0.099$). The second is a muted average quantity response (-0.846). The third is heterogeneous quantity reallocation: the cost shock reduces quantities on cash relationships ($\hat{\beta}_1 = -1.492$) while the interaction with TC intensity is large and positive ($\hat{\beta}_3 = 1.475$). The model approximates the price response as uniform for tractability; the state-dependent pricing patterns documented in Table 13 are discussed as a natural extension.

The model is intentionally parsimonious. Its purpose is to formalize one interpretation of the empirical patterns, map the estimated quantity coefficients into structural parameters, and discipline a counterfactual quantifying the cost of removing priority-based allocation. In the model, trade-credit intensity maps monotonically into relational surplus, providing a theoretical foundation for interpreting the empirical proxy. Section 8 outlines how this framework can be enriched into a fully dynamic, general-equilibrium production-network model.

6.1 Environment

Agents and timing. Consider a single supplier j that interacts repeatedly with a unit mass of downstream buyers over an infinite horizon, with common discount factor $\beta \in (0, 1)$, where β measures the weight the supplier places on future profits relative to current profits. The model focuses on a single supplier in isolation in order to derive closed-form predictions; this is a partial-equilibrium simplification rather than an assumption that the supplier is a monopolist in its product market. In the data, suppliers compete with other firms selling similar products, and the competitive constraints discussed below reflect precisely this discipline. In each period t, the supplier draws an i.i.d. cost shock ν_t with $\mathbb{E}[\nu] = 0$ and bounded support, observes the realization, and then sets prices and buyer-specific quantities $\{q_{it}\}$.

Buyer types and relational status. Buyers differ in their pre-existing relational status with the supplier. Trade-credit (TC) buyers, of measure $\mu^{TC} \in (0, 1)$, represent long-standing relationships in which the supplier extends trade credit and provides non-contractible relationship-specific services such as reliable delivery timing, flexible order sizes, and after-sale support. Cash buyers, of measure $\mu^{C} = 1 - \mu^{TC}$, represent arm's-length, spot-market transactions with no trade credit and no relationship-specific surplus. The parameter μ^{TC} governs the share of the supplier's buyer base that is relational. I take this partition as given,

reflecting the pre-shock relational structure observed in the data. Endogenizing relationship formation is discussed in Section 8.

Demand. Each buyer i has iso-elastic demand augmented by a relationship-specific demand shifter:

$$q_i(p, s_i) \;=\; A(s_i)\, p^{-\epsilon}, \qquad \epsilon > 1, \tag{2}$$

where $s_i \geq 0$ is the level of non-contractible "relational service" provided by the supplier, with $A(0) = 1$, $A'(s) > 0$, and $A''(s) \leq 0$. The parameter ϵ is the price elasticity of demand: it governs how responsive buyers are to changes in the supplier's price, with higher ϵ implying greater price sensitivity and less supplier pricing power. The term "iso-elastic" means that this elasticity is constant—a 1% price increase always reduces demand by ϵ%, regardless of the price level. The restriction $\epsilon > 1$ is a standard regularity condition ensuring that the supplier's revenue is well-behaved at the optimum. Trade-credit buyers receive $s^{TC} > 0$, so their demand is shifted outward by the factor $A(s^{TC}) > 1$; cash buyers operate at the spot baseline $s^C = 0$, so that $A(0) = 1$.

This formulation captures in reduced form the idea, documented empirically by Garcia-Marin et al. (2023) and modeled theoretically by Cuñat (2007), that trade credit is bundled with relationship-specific services that raise the buyer's willingness to purchase from the supplier. The demand premium $A(s^{TC}) - 1$ generates the relational surplus that the supplier has an incentive to protect.

Production technology. The supplier produces total output $Q_t = \int q_{it}\, di$ at cost

$$C(Q; \nu) \;=\; (c + \nu)\, Q \;+\; \frac{\gamma}{2}\, Q^2, \qquad \gamma > 0, \tag{3}$$

where $c > 0$ is the baseline marginal input cost and $\gamma > 0$ is the cost convexity parameter. The first term, $(c+\nu)Q$, represents the linear component of variable costs: each unit of output requires $c + \nu$ in variable inputs, where ν is the realized cost shock. The second term, $\frac{\gamma}{2}Q^2$, introduces a quadratic component that makes marginal cost—equal to $c+\nu+\gamma Q$—increasing in total output. This convexity reflects short-run capacity constraints that are relevant for the suppliers in the sample. As a firm scales up production within a period, it may need to operate machinery at higher intensity, schedule additional labor shifts at premium wages (i.e., the additional compensation paid to workers for hours beyond the standard schedule), source marginal units of imported intermediates from more costly channels, or stretch its logistics and distribution capacity. The parameter γ governs the severity

of these short-run constraints: a higher γ means marginal costs rise more steeply with output.[5]

The convexity of $C(\cdot)$ is essential for generating quantity rationing across buyers. With constant marginal costs ($\gamma = 0$), the supplier would be indifferent about which buyers' quantities to adjust, because the cost saving from reducing one buyer is the same as from reducing another. With convex costs ($\gamma > 0$), the supplier's marginal cost depends on total output, so reallocating a unit of output from one buyer to another changes the cost of serving the entire buyer base. This interdependence is what creates a meaningful allocation problem: the supplier must decide how to distribute a given level of output across heterogeneous buyers, even when the aggregate output response is small.

6.2 Relational contracting and the incentive-compatibility constraint

The relational service s^{TC} is non-contractible: no court can verify delivery reliability or credit extension. The supplier's willingness to maintain these services must therefore be sustained by self-enforcing incentives, as in the relational-contracting frameworks of Baker et al. (2002) and Levin (2003).

Stage game and deviation. In each period, after observing ν_t, the supplier can deviate on any TC buyer by withholding relational service and reducing the buyer's quantity below the cooperative level. If the supplier deviates, the buyer observes the deviation (a short shipment, a refusal to extend credit), and the relationship reverts permanently to the cash regime: $s_i = 0$ and continuation value zero.[6]

Continuation values. Let V^{TC} denote the supplier's expected per-period profit from a single TC relationship in the cooperative equilibrium and V^{C} the corresponding profit from a cash relationship. Because TC buyers have higher demand at the same price, $V^{TC} > V^{C}$. The supplier's continuation value from maintaining a TC relationship is

$$J^{TC} \;=\; \frac{\beta}{1-\beta}\big(V^{TC} - V^{C}\big), \tag{4}$$

5. Since the cost shock ν_{jt} is constructed from the firm's lagged exposure to imported intermediates interacted with the log change in the exchange rate, γ can be interpreted as the rate at which the shadow cost of imported inputs rises with the scale of production during episodes of peso depreciation.

6. The grim-trigger punishment assumption is standard in the relational-contracting literature and can be relaxed to renegotiation-proof punishment schemes without changing the qualitative results (Levin 2003).

representing the discounted stream of relational quasi-rents that would be forfeited upon deviation.

Incentive-compatibility constraint. In any period with cost realization ν, denote by $\pi^{dev}(\nu)$ the one-period profit from optimally deviating on a single TC buyer (treating it as cash for one period) and by $\pi^{coop}(\nu)$ the cooperative profit. The relational contract is self-enforcing if and only if, for every realization of ν:

$$\underbrace{\pi^{dev}(\nu) - \pi^{coop}(\nu)}_{\text{temptation to deviate}} \leq \underbrace{\beta\, J^{TC}}_{\text{value of preserving the relationship}} . \tag{5}$$

The left-hand side is the temptation to deviate—the short-run profit gain from reducing a TC buyer's quantity. The right-hand side is the discounted continuation value of maintaining the relationship. The constraint is tightest when ν is large and positive, because that is when the supplier most wants to reduce production and the temptation to ration TC buyers is greatest.

Reduced-form quantity floor. When (5) holds, the supplier's optimal policy keeps TC quantities at least at the cooperative level that buyers expect. For closed-form comparative statics, I summarize the binding IC condition as

$$q_t^{TC} \;\geq\; \bar{q}, \tag{6}$$

where $\bar{q}$ is the cooperative steady-state quantity (derived below). This is not an ad hoc assumption but a direct consequence of (5): provided J^{TC} is sufficiently large, the supplier strictly prefers maintaining $q^{TC} \geq \bar{q}$ to deviating and losing the relationship.[7]

Remark 1 (Mapping $\bar{q}$ to trade-credit intensity) *In the data, relational priority is proxied by* $\mathrm{TC}^{\mathrm{pre}}_{ij,t}$. *In the model, $\bar{q}$ is pinned down by the relational surplus J^{TC}, which depends on s^{TC} and the demand premium $A(s^{TC})$. The mapping is monotone: higher-surplus relationships have higher $\bar{q}$ and, empirically, higher trade-credit shares. The model therefore provides a theoretical foundation for using* $\mathrm{TC}^{\mathrm{pre}}_{ij,t}$ *as a proxy for relational priority in the supplier's allocation decision: higher-surplus relationships map into higher trade-credit shares.*

7. More precisely, $\bar{q}$ is the maximum quantity floor consistent with (5) holding at the worst-case shock. If shocks can be very large, $\bar{q}$ must be set lower. I calibrate $\bar{q}$ to the steady-state level, which is appropriate as long as observed shocks are moderate relative to J^{TC}.

6.3 Pricing constraints and the quantity allocation

The estimation results show that, in the pooled specification, the average price interaction with TC intensity is near zero ($\hat{\beta}_3 = -0.093$), but that under adverse shocks specifically, TC buyers face somewhat higher prices ($\hat{\beta}_3 = 0.731$ on price). The model focuses on the quantity allocation—which is where the strongest and most robust heterogeneity appears—and approximates the price as fixed at p^* for tractability. This approximation is motivated by three considerations.

First, the shocks are idiosyncratic to the supplier. Because the shock raises costs for one supplier but not for its competitors, the supplier's ability to pass through the cost increase is limited by competitive pressure, consistent with incomplete pass-through as documented in the ERPT literature (Amiti et al. 2014; Berman et al. 2012). Cash buyers, who transact at arm's length, face particularly strong competitive alternatives; TC buyers, who may face higher switching costs due to the value of relational services, may accept modest price increases. This asymmetry is consistent with the adverse-shock pricing pattern in Table 13.

Second, even when prices adjust differentially under adverse shocks, the quantity insulation of TC links is preserved and even strengthened, because the demand-side effect of higher TC prices would push TC quantities *down*. The supply-side quantity protection must therefore be large enough to offset both the aggregate output contraction and the demand reduction from the price increase. The quantity allocation mechanism, which is the central object of the model, operates regardless of the price pattern.

Third, under iso-elastic demand, a supplier with market power would adjust prices by an amount of order $\frac{1}{\epsilon-1}$, which is small for large ϵ. Many suppliers in the sample operate in product markets with close substitutes, making the fixed-price approximation quantitatively reasonable.

For tractability, the model therefore holds $p \approx p^*$ and focuses on the quantity allocation problem. Section 8.1 discusses how endogenizing buyer-type-specific pricing would generate the state-dependent pricing patterns documented in Table 13 while preserving the central quantity-reallocation prediction.

6.4 Why different margins for different buyers

The empirical finding that the strongest heterogeneity is in quantities, with prices showing a near-zero average interaction but state-dependent differential pricing under adverse shocks, reflects the fact that the supplier faces different constraints over different buyer types.

Cash buyers transact at arm's length and can likely source from competitors at the unchanged market price. Raising the cash-buyer price risks losing the buyer to competitors whose costs have not risen. But cash buyers have no relational contract guaranteeing a specific quantity, and cutting their shipments incurs no future penalty. The adjustment margin for cash buyers is therefore quantities.

TC buyers may face higher switching costs due to the value of relational services—credit terms, reliable delivery, after-sale support—that are costly to replicate with a new supplier. The supplier may therefore be able to raise prices for these buyers without losing them (consistent with the positive price interaction under adverse shocks). But short-shipping a TC buyer signals unreliability and threatens the continuation value J^{TC} that sustains the relationship. The adjustment margin for TC buyers is therefore prices rather than quantities.

When a cost shock hits, the supplier adjusts on the margin where it has leverage with each buyer type. It maintains TC quantities (preserving the relationship) while potentially raising TC prices modestly. It holds cash-buyer prices approximately constant (competitive alternatives are unchanged) while rationing their quantity (no relational penalty). The model formalizes the quantity allocation—the central and most robust margin—by approximating the price as fixed at p^* and focusing on the quantity floor $q^{TC} \geq \bar{q}$ and the residual allocation to cash buyers.

6.5 Equilibrium analysis

6.5.1 Steady state ($\nu = 0$)

In the absence of shocks, the supplier sets the price p^* to maximize static profits over the full buyer base. Total demand at price p is

$$Q^D(p) \;=\; \left[\mu^{TC} A(s^{TC}) + \mu^C\right] p^{-\epsilon} \;\equiv\; \bar{A}\, p^{-\epsilon}, \tag{7}$$

where $\bar{A} = \mu^{TC} A(s^{TC}) + \mu^{C}$ is a demand-weighted average across buyer types. The first-order condition for the optimal price yields

$$p^* \;=\; \frac{\epsilon}{\epsilon - 1}\big(c + \gamma Q^*\big), \tag{8}$$

where $Q^* = \bar{A}\,(p^*)^{-\epsilon}$ is steady-state total output. The ratio $\frac{\epsilon}{\epsilon-1}$ is the standard markup over marginal cost in models of monopolistic competition: it is decreasing in ϵ, so that as buyers become more price-sensitive the supplier's markup declines. Individual steady-state quantities are

$$q^{TC,*} = A(s^{TC})\,(p^*)^{-\epsilon}, \tag{9}$$

$$q^{C,*} = (p^*)^{-\epsilon}. \tag{10}$$

Because $A(s^{TC}) > 1$, it follows that $q^{TC,*} > q^{C,*}$: trade-credit relationships transact higher volumes in steady state. The relational threshold is set at $\bar{q} = q^{TC,*}$, and the per-period relational surplus is $V^{TC} - V^{C} = (p^* - c - \gamma Q^*)(q^{TC,*} - q^{C,*}) > 0$.

6.5.2 Short-run response to a cost shock ($\nu > 0$)

Suppose the supplier draws $\nu > 0$. The price remains at p^*. The supplier determines total output by equating the fixed price to the new short-run marginal cost:

$$p^* \;=\; c + \nu + \gamma\, Q(\nu) \qquad \Longrightarrow \qquad Q(\nu) \;=\; \frac{p^* - c - \nu}{\gamma}. \tag{11}$$

The aggregate contraction relative to steady state is

$$\Delta Q \;\equiv\; Q(\nu) - Q^* \;=\; -\frac{\nu}{\gamma}. \tag{12}$$

The allocation problem. The supplier chooses (q^{TC}, q^{C}) to maximize

$$\Pi(\nu) \;=\; p^*\, Q - (c + \nu)\, Q - \frac{\gamma}{2}\, Q^2 \;+\; \beta\, \mu^{TC}\, J^{TC}\, \mathbf{1}\{q^{TC} \geq \bar{q}\}, \tag{13}$$

where $Q = \mu^{TC} q^{TC} + \mu^{C} q^{C}$, subject to $q^{TC} \geq 0$ and $q^{C} \geq 0$. The last term is the discounted value of preserving all TC relationships.

The supplier compares the consequences of reducing each type of buyer's quantity. Reducing q^{C} lowers current revenue but incurs no dynamic penalty. Reducing q^{TC} below $\bar{q}$ triggers the discrete loss of $\beta \mu^{TC} J^{TC}$. Provided J^{TC} is sufficiently large, the supplier protects TC

links.

Solution. Because the IC constraint (5) holds, the supplier sets $q^{TC} = \bar{q} = q^{TC,*}$ and assigns the residual output to cash buyers:

$$q^{TC}(\nu) = q^{TC,*} \qquad \text{(TC links fully insured)}, \tag{14}$$

$$q^{C}(\nu) = \frac{Q(\nu) - \mu^{TC}\, q^{TC,*}}{\mu^{C}} \;=\; q^{C,*} - \frac{\nu}{\mu^{C}\gamma} \qquad \text{(cash absorbs the full adjustment)}. \tag{15}$$

This interior solution holds for $\nu \in [0, \bar{\nu}]$, where $\bar{\nu} = \mu^{C}\gamma\, q^{C,*}$ is the shock at which cash quantities reach zero.

6.6 Comparative statics and mapping to the estimation results

Proposition 1 (Differential pass-through of cost shocks) *Under convex costs, relational priority, and fixed-price approximation, a cost shock $\nu > 0$ generates: (1) $\frac{\partial q^{TC}}{\partial \nu} = 0$, so TC quantities do not respond; (2) $\frac{\partial q^{C}}{\partial \nu} = -\frac{1}{\mu^{C}\gamma} < 0$, so cash quantities decline, with the response amplified by the factor $1/\mu^{C}$; (3) $p \approx p^{*}$ (tractability approximation; the data show small price adjustments that do not alter the quantity allocation); and (4) the supplier's return on assets falls and leverage rises.*

Proof. Parts 1–2 follow from differentiating (14)–(15). Part 3 is the maintained approximation: the pooled data show a near-zero average price differential ($\hat{\beta}_3 = -0.093$); the model approximates the price as $p \approx p^{*}$, which is appropriate for the quantity allocation problem. For Part 4, at approximately fixed p^{*} the cost increase raises total costs while total revenue falls (because Q declines); net income therefore falls, and if the supplier borrows to finance the maintained TC deliveries at higher input costs, liabilities rise relative to assets. ■

Mapping to the main link-level estimates. The empirical specification estimates $y_{ijbt} = \beta_1\, \nu_{jt} + \beta_2\, \mathrm{TC}^{\mathrm{pre}}_{ij,t} + \beta_3\, \nu_{jt} \times \mathrm{TC}^{\mathrm{pre}}_{ij,t} + \mathrm{FE} + \varepsilon_{ijbt}$. In the model, β_1 captures the response of pure cash links ($\mathrm{TC}^{\mathrm{pre}} = 0$) and $\beta_1 + \beta_3 \cdot \mathrm{TC}^{\mathrm{pre}}$ the response at a given TC share. The model predicts $\beta_1 < 0$ (cash links contract; the estimate is $\hat{\beta}_1 = -1.492$), $\beta_3 > 0$ (high-TC links are insulated; the estimate is $\hat{\beta}_3 = 1.475$), and a near-zero average price interaction ($\hat{\beta}_3 = -0.093$ on price), with state-dependent differential pricing under adverse shocks ($\hat{\beta}_3 = 0.731$ on price in Table 13). At the median pre-shock TC share of approximately 0.80, the predicted net quantity effect is $\hat{\beta}_1 + \hat{\beta}_3 \times 0.80 \approx -0.31$, a modest decline that is much smaller than the -1.49 decline on pure cash links.

The point estimate at $\text{TC}^{\text{pre}} = 1$. Extrapolating to $\text{TC}^{\text{pre}} = 1$ yields a net quantity effect of $\hat{\beta}_1 + \hat{\beta}_3 \approx -0.02$, essentially zero. The baseline model predicts exactly zero. Very few relationships in the data have TC^{pre} close to one; at the median high-TC share the estimates are consistent with the model's prediction of approximately zero. A mild positive effect at the extreme could be rationalized if the relational constraint (6) is slack for the highest-surplus relationships, so that the supplier actively reallocates marginal units toward its most valuable partners when capacity is constrained. This extension is discussed in Section 8.4.

Mapping to the supplier-level financial regressions. A positive cost shock lowers net income at roughly constant revenue (since TC links, which account for the bulk of volume, are maintained), reducing ROA. To finance maintained deliveries at higher input costs, the supplier may increase its reliance on external finance, raising leverage. The estimated coefficients ($\hat{\theta}^{ROA} = -0.111$; $\hat{\theta}^{Lev} = 0.099$) are consistent with this prediction. The profit margin coefficient is negative (-0.025) but statistically insignificant, consistent with the margin compression being partially offset by a compositional shift in the supplier's sales mix: as cash volumes contract and TC volumes are maintained, the average unit sold carries a higher relational markup, attenuating the decline in the net-income-to-sales ratio.[8]

Mapping to the average-effects specification. The average quantity response across all matches is a weighted average of the TC response (zero) and the cash response (negative): $\frac{\partial \bar{q}}{\partial \nu} = \mu^{TC} \cdot 0 + \mu^{C} \cdot (-\frac{1}{\mu^{C}\gamma}) = -\frac{1}{\gamma}$. With sufficiently high γ (steep marginal costs), this is small, consistent with the insignificant point estimate of -0.846. The average price response is approximately zero in the model (fixed-price approximation); the data show a small but significant average increase (0.440), confirming that the approximation is reasonable and that the price margin does not drive the quantity reallocation.

Amplification through concentration. The factor $1/\mu^{C}$ in $\partial q^{C}/\partial \nu$ reflects an amplification mechanism: because TC buyers are effectively senior claimants on the supplier's scarce capacity, the entire aggregate output reduction is concentrated on the residual cash segment. With $\mu^{C} \approx 0.20$–0.30, each cash buyer absorbs a quantity reduction roughly 3–5 times larger than the aggregate per-buyer contraction.

8. In the model, total revenue is $p^{*}Q(\nu)$, which falls, but revenue from TC buyers is unchanged. The profit margin $\Pi(\nu)/(p^{*}Q(\nu))$ has a falling numerator but also a falling denominator, partially offsetting the decline.

6.7 Profitability absorption and the cost of priority-based allocation

The mechanism has a direct cost to the supplier. In the shock state, total profit is

$$\Pi(\nu) = p^* Q(\nu) - (c + \nu)Q(\nu) - \frac{\gamma}{2}Q(\nu)^2. \tag{16}$$

Because the price is fixed and total output falls only partially (TC quantities are maintained), the supplier's revenue declines by less than the increase in total costs. The gap is absorbed in lower margins. Return on assets declines because net income falls while assets do not shrink proportionally. Leverage increases if the supplier borrows to cover higher input costs while maintaining deliveries to TC buyers. This is the cost of priority-based allocation, and the supplier bears it because the continuation value J^{TC} exceeds the one-period profit sacrifice, as guaranteed by (5).

6.8 When does the relational contract break down?

The model identifies a threshold $\bar{\nu} = \mu^C \gamma \, q^{C,*}$ beyond which priority-based allocation fails. If $\nu > \bar{\nu}$, the supplier has already reduced cash quantities to zero and must begin reducing TC quantities; at that point the temptation to deviate may exceed the continuation value, and the relational contract unravels.

This threshold connects to the broader literature. Bocola and Bornstein (2025) show that trade credit amplifies output declines during severe crises, precisely because the relational contracts sustaining trade credit break down when shocks are sufficiently large. The model predicts that idiosyncratic cost shocks within the normal range—such as the exchange-rate-driven shocks studied here—are absorbed by the relational network, but that aggregate or very large shocks can overwhelm the mechanism and trigger a shift from dampening to amplification.

7 Quantitative Implications

The model maps the estimated regression coefficients into structural parameters and a counterfactual. Even in its parsimonious form, this mapping delivers two results that go beyond rationalizing the reduced-form estimates.

The first exercise disciplines the short-run convexity of supplier costs—a parameter

that governs how easily firms can scale production and that is central to any quantitative production-network model, yet is rarely pinned down from micro data. The key insight is that the cross-sectional heterogeneity in pass-through (the cash-vs.-TC differential) maps into a supply-side technology parameter: the supplier's tendency to reduce cash quantities rather than raise prices or absorb the full shock in margins reveals how steeply its marginal costs rise with output.

The second exercise asks what downstream quantity volatility would look like if priority-based allocation did not exist. The standard production-network literature (Acemoglu et al. 2012; Gabaix 2011; Bigio and La'O 2020; Altinoglu 2021) typically abstracts from heterogeneity across relationships within a supplier's buyer base. The estimates show this abstraction misses a first-order margin. The same shock originating at the same supplier generates a 3 percentage-point gap in quantity responses across relationship types, depending entirely on pre-existing payment terms.

7.1 Recovering the cost convexity parameter (γ)

Short-run cost convexity is a key input for production-network models: it determines whether cost shocks are absorbed by suppliers in lower margins, passed through to buyers in higher prices, or accommodated through quantity rationing. Despite its importance, γ is difficult to estimate directly because researchers rarely observe the supplier's full cost function. The priority-based allocation model offers an indirect identification strategy: because the supplier protects TC links and concentrates all quantity adjustment on cash links, the observed cash-buyer elasticity $\hat{\beta}_1$ disciplines γ given observable moments of the supplier's customer base.

From Proposition 1, the semi-elasticity of cash-buyer quantities with respect to the cost shock is

$$\frac{\partial \ln q^C}{\partial \nu} = \frac{1}{q^{C,*}} \frac{\partial q^C}{\partial \nu} = -\frac{1}{\mu^C \gamma \, q^{C,*}}. \tag{17}$$

Mapping the regression coefficient to (17):

$$\hat{\beta}_1 \;\approx\; -\frac{c}{\mu^C \gamma \, q^{C,*}}. \tag{18}$$

Solving for the implied convexity:

$$\gamma_{\text{implied}} = -\frac{1}{\hat{\beta}_1} \frac{c}{\mu^C \, q^{C,*}}. \tag{19}$$

To compute γ_{implied}, I use $\hat{\beta}_1 \approx -1.5$ from Table 8, $\mu^C \approx 0.2$–0.3 from the sample distribution of $\text{TC}_{ij,t}^{\text{pre}}$, and $q^{C,*}$ and c recovered from sample averages of per-link trade volumes and the supplier's cost of goods sold, respectively. The implied ratio $\gamma Q^*/c \approx 2.4$ indicates significant short-run cost convexity: steady-state marginal cost is roughly 3.4 times the baseline linear component, implying that a 10% increase in output raises marginal cost by approximately 7%.

The broader point is methodological. Standard approaches to estimating cost convexity require observing output and input prices at the plant level, which is feasible only in manufacturing census data for a limited set of countries. The priority-based allocation framework offers an alternative: if one observes heterogeneous quantity responses across relationship types for the same supplier, the differential response identifies γ without requiring direct observation of the supplier's input prices or production function.

7.2 Counterfactual: the economy without priority-based allocation

The estimates show that a one-standard-deviation cost shock generates a 3 percentage-point gap in quantity responses between cash and TC buyers of the same supplier. This within-supplier heterogeneity is not captured by models that typically abstract from heterogeneity across relationships within a supplier's buyer base (Acemoglu et al. 2012; Gabaix 2011; Bigio and La'O 2020; Altinoglu 2021). To quantify how much this reallocation channel matters, I ask: what would the economy look like if relational enforcement failed and all relationships reverted to arm's-length, cash-based trade?

Connection to the existing literature. A central finding of the recent trade-credit literature is that trade credit plays a dual role in production networks: it relaxes financial constraints and raises economic activity in normal times, but creates fragility and can amplify downturns in adverse times. Bocola and Bornstein (2025) document this at the aggregate level. Altinoglu (2021) finds that adding trade-credit linkages to a production-network model substantially increases the share of aggregate volatility explained by financial shocks. On the stabilizing side, Cazzaniga et al. (2025) show that trade-credit-providing suppliers reduce output losses for downstream firms facing financial frictions, and Ersahin et al. (2024) document that trade-credit flows increase when firms face operating shocks.

The counterfactual exercise connects to this macro-level dual role by providing its micro-foundation. The existing literature identifies the dual role through the lens of aggregate financial constraints and credit supply. This paper shows that a distinct mechanism—relational priority in the supplier's quantity allocation—generates both a level benefit and a

volatility redistribution at the individual link level.

Setup. I consider a counterfactual in which relational enforcement fails ($J^{TC} = 0$). All relationships revert to the spot regime: the supplier no longer provides relational services ($s_i = 0$ for all i), and the demand premium that sustained high TC volumes vanishes. The counterfactual operates through two margins.

Margin 1: Steady-state output loss. In the baseline, steady-state total output is $Q^* = \bar{A}\,(p^*)^{-\epsilon}$, where $\bar{A} = \mu^{TC} A(s^{TC}) + \mu^C$. In the counterfactual, all buyers have $A(s_i) = 1$, so total output falls. Holding the price constant:

$$\frac{Q^*_{cf}}{Q^*} = \frac{1}{\bar{A}} = \frac{1}{\mu^{TC} A(s^{TC}) + \mu^C} < 1. \tag{20}$$

Using $\mu^{TC} = 0.7$ and $A(s^{TC}) = 1.24$ (the demand premium implied by the volume advantage of TC relationships), steady-state output falls by approximately 15%. This is a permanent loss driven entirely by the destruction of relationship capital.

This level effect is the micro-level counterpart of the normal-times benefit documented in the macro literature. Bocola and Bornstein (2025) show that trade credit relaxes borrowing constraints and supports aggregate output; here, the mechanism operates through the relational demand premium—the services, reliability, and flexible terms documented by Garcia-Marin et al. (2023) and Benguria et al. (2023)—that would disappear under arm's-length trade.

Margin 2: Volatility redistribution. Conditional on the lower steady-state output, the supplier spreads the aggregate contraction proportionally across all buyers:

$$\Delta \ln q^{TC}_{\text{cf}} = \Delta \ln q^{C}_{\text{cf}} = \frac{\Delta Q}{Q^*_{cf}} = -\frac{\nu}{\gamma Q^*_{cf}}. \tag{21}$$

In the data, a one-standard-deviation cost shock reduces cash-buyer quantities by approximately 3% while leaving TC-buyer quantities unchanged. The aggregate output contraction is approximately $\mu^C \times 3\% \approx 0.9\%$. In the counterfactual, this 0.9% contraction would be spread evenly across all buyers. TC buyers would therefore go from zero quantity volatility to a 0.9% decline per one-standard-deviation shock episode.

This volatility margin connects to the broader finding that relational structures reshape

the distribution of risk. Cazzaniga et al. (2025) show that trade credit stabilizes downstream firms facing financial frictions, but from the demand side and without identifying who bears the displaced risk. The counterfactual here provides the supply-side complement.

Relief for cash buyers. Cash buyers would experience lower volatility in the counterfactual: -0.9% instead of -3% per one-standard-deviation shock. Priority-based allocation raises cash-buyer quantity volatility by roughly 2 percentage points per shock episode, concentrating the full adjustment on about 30% of the supplier's customer base. However, this volatility reduction comes at the cost of permanently lower steady-state quantities.

Caveats. This counterfactual is partial equilibrium and holds buyer behavior fixed. If priority-based allocation disappeared, buyers would adjust their sourcing strategies, so the counterfactual provides an upper bound. The partition into TC and cash buyers is taken as given, and in a dynamic model relationship formation would respond endogenously. The exercise is best interpreted as the short-run effect of a sudden loss of relational enforcement.

Summary. The macro literature has established that trade credit plays a dual role in production networks: supporting output in normal times and amplifying downturns in crises (Bocola and Bornstein 2025; Altinoglu 2021). The counterfactual provides the micro-foundation for this duality. At the link level, relational priority simultaneously raises the level of economic activity—by sustaining high-demand, high-surplus relationships—and reshapes the distribution of volatility—by insulating those relationships from upstream cost shocks at the expense of spot buyers. Altinoglu (2021) shows that trade credit raises the aggregate volatility contribution of financial shocks by one-third; the estimates here identify where that volatility lands. Bocola and Bornstein (2025) show that the amplification margin activates in crises; the estimates here show that a within-supplier amplification operates continuously through routine idiosyncratic shocks. For macroeconomic models of shock propagation through production networks (Acemoglu et al. 2012; Gabaix 2011; Bigio and La'O 2020), the implication is that abstracting from within-supplier heterogeneity in buyer relationships misses both a level and a volatility margin that are quantitatively important.

7.3 Implications for efficiency and reallocation

The reallocation mechanism has ambiguous welfare implications. On one hand, by insulating high-volume relationships, the mechanism protects matches with the highest relational surplus, preserving the core of the production network during temporary shocks. On the other hand, if cash buyers are young, high-growth entrants who have not yet accumulated relationship

capital, the mechanism may act as a barrier to entry: by concentrating aggregate volatility onto unconnected firms, the relational network may dampen their growth, raise their exit risk, and distort long-run industry dynamics in favor of incumbents. Distinguishing between these interpretations requires data on cash-buyer characteristics and life-cycle dynamics, which is part of the ongoing research agenda.

8 Discussion and Extensions

The model above formalizes and disciplines the empirical findings with minimal structural assumptions. Several extensions are needed to build a framework suitable for quantitative production-network analysis.

8.1 Endogenous price adjustment

The current model approximates the price as fixed at p^*. A natural extension introduces buyer-type-specific pricing subject to two constraints: a competitive constraint $p^C(\nu) \leq \bar{p}$ (cash buyers likely have competitive alternatives) and a relational quantity floor $q^{TC}(\nu) \geq \bar{q}$ (TC buyers must receive their committed quantity). Under adverse shocks, this would generate the pattern documented in Table 13: the supplier raises p^{TC} (TC buyers may face higher switching costs), holds $p^C \approx \bar{p}$ (cash buyers have competitive alternatives), maintains $q^{TC} = \bar{q}$ (relational commitment), and reduces q^C (residual allocation). The key quantity prediction is preserved and strengthened: even when TC buyers face higher prices, their quantities are maintained, implying that the supply-side protection must offset both the output contraction and the demand-side effect of the price increase. Jointly estimating the price and quantity responses could recover ϵ and the relational pricing wedge as additional structural parameters.

8.2 Dynamics and relationship formation

The static allocation problem abstracts from intertemporal smoothing and the endogenous evolution of relationship capital. A fully dynamic model would introduce persistent shocks, under which the IC constraint depends on the expected future path and $\bar{q}$ becomes state-dependent. It would also incorporate relationship-specific investment, with the continuation value J^{TC} evolving endogenously with relationship age and history. Finally, it would allow for entry, exit, and switching: cash buyers who are repeatedly rationed may exit or switch to alternative suppliers, while new buyers enter and gradually build relationships. The partial-equilibrium allocation rule derived here provides the micro-foundation for each supplier node,

which can then be aggregated across the production network.

8.3 General equilibrium and network effects

In the current framework, each supplier is analyzed in isolation. In a production network, the quantity reductions imposed on cash buyers by one supplier become input shocks for those buyers, potentially cascading downstream. A general-equilibrium model would trace these cascading effects and assess whether priority-based allocation at the bilateral level translates into aggregate stability or merely redistributes volatility. This is related to Altinoglu (2021) and Bocola and Bornstein (2025), who show that trade-credit linkages can either dampen or amplify aggregate volatility depending on network structure and the distribution of financial constraints.

8.4 Flight to quality and slack relational constraints

The baseline model predicts exactly zero quantity response for TC links. The main estimates show that at the median TC share the net response is approximately -0.31, a modest decline, while at $\text{TC}^{\text{pre}} = 1$ the net effect is $\hat{\beta}_1 + \hat{\beta}_3 \approx -0.02$, essentially zero. If the relational constraint (6) is slack for the highest-surplus relationships, the supplier may actively shift marginal units toward its most valued partners when capacity is constrained. Formally, replacing the hard floor with a smooth relational-value function would yield an interior optimum in which $q^{TC}(\nu) > q^{TC,*}$ for small shocks.

8.5 Asymmetric shocks and two-sided allocation

The model can also accommodate demand-side shocks hitting downstream buyers. In that case, the direction of priority reverses: a TC buyer hit by a negative demand shock reduces purchases, and the supplier must decide whether to accommodate the reduction or enforce minimum commitments. Cazzaniga et al. (2025) study this direction and find that trade-credit-providing suppliers stabilize downstream firms, consistent with two-sided priority-based allocation. Unifying both directions would provide a complete characterization of how the relational structure of payment terms mediates shock transmission from both upstream and downstream.

These extensions constitute the next phase of this research agenda. The key insight from the current paper—that the dimension of relationship heterogeneity captured by pre-shock trade-credit intensity determines the cross-sectional allocation of cost shocks across buyer–supplier links—provides the empirical discipline that any richer model must match.

9 Conclusion

This paper shows that upstream input-cost shocks at suppliers are not transmitted uniformly across their buyer base. Exploiting within seller–buyer–product match variation over time, the estimates show that the strongest heterogeneity in pass-through is in quantities: when a supplier faces a cost increase, quantities in its cash relationships decline substantially while quantities in its high-trade-credit relationships remain largely unchanged. Under adverse shocks, the supplier also raises prices for TC buyers—consistent with TC buyers facing higher switching costs—while rationing cash buyers on the quantity margin. At the supplier level, the shocks reduce return on assets and increase leverage, indicating that suppliers absorb part of the cost increase in their margins and balance sheets.

Pre-shock trade-credit intensity serves as a predetermined summary statistic of relationship type that predicts how suppliers allocate cost shocks across their buyer base. The TC interaction survives a saturated specification controlling for buyer and supplier size, age, export status, and relationship length; relationship length does not predict the differential response; and the state-dependent pricing pattern—in which TC buyers face higher prices under adverse shocks but not under favorable shocks—is consistent with the relational-contracting interpretation in which the supplier uses different adjustment margins for different buyer types.

A partial-equilibrium model with convex costs and relational contracting provides one micro-foundation for these patterns, maps the estimated quantity elasticities into structural parameters, and delivers a counterfactual showing that priority-based allocation simultaneously raises the level of interfirm trade (by sustaining relationship capital) and reshapes the distribution of upstream cost risk (by insulating high-trade-credit partners at the expense of cash buyers). From a policy perspective, the results imply that the structure of payment terms within supplier–buyer networks is a key determinant of which firms bear the real consequences of upstream cost shocks. Embedding the estimated heterogeneous pass-through patterns into a quantitative production-network model is a natural next step to assess aggregate implications for volatility, misallocation, and welfare.

References

Acemoglu, D., V. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica.*

Altinoglu, L. 2021. "The origins of aggregate fluctuations in a credit network economy." *Journal of Monetary Economics.*

Amiti, M., O. Itskhoki, and J. Konings. 2014. "Importers, Exporters, and Exchange Rate Disconnect." *American Economic Review* 104 (7): 1942–1978.

Baker, G., R. Gibbons, and K. Murphy. 2002. "Relational Contracts and the Theory of the Firm." *The Quarterly Journal of Economics.*

Benguria, F., A. Garcia-Marin, and T. Schmidt-Eisenlohr. 2023. "Trade Credit and Relationships." *Working Paper.*

Berman, N., P. Martin, and T. Mayer. 2012. "How do Different Exporters React to Exchange Rate Changes?" *The Quarterly Journal of Economics* 127 (1): 437–492.

Bigio, S., and J. La'O. 2020. "Distortions in Production Networks." *The Quarterly Journal of Economics.*

Bocola, L., and G. Bornstein. 2025. "The Macroeconomics of Trade Credit." *Working Paper NBER.*

Cazzaniga, M., P. Pannella, and L. Alencar. 2025. "Trade Credit in a Developing Country: The Role of Large Suppliers in the Production Network." *Working Paper.*

Chodorow-Reich, G. 2014. "The Employment Effects of Credit Market Disruptions: Firm-level Evidence from the 2008–9 Financial Crisis." *The Quarterly Journal of Economics.*

Costello, A. 2020. "Credit Market Disruptions and Liquidity Spillover Effects in the Supply Chain." *Journal of Political Economy.*

Cuñat, V. 2007. "Trade credit: suppliers as debt collectors and insurance providers." *The Review of Financial Studies* 2 (20): 491–527.

Ersahin, N., M. Giannetti, and R. Huang. 2024. "Trade credit and the stability of supply chains." *Journal of Financial Economics.*

———. 2025. "Supply chain risk: Changes in supplier composition and vertical integration." *Journal of International Economics.*

Gabaix, X. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica.*

Garcia-Marin, A., S. Justel, and T. Schmidt-Eisenlohr. 2023. "Trade Credit, markups and Relationships." *Working Paper.*

Gopinath, G., O. Itskhoki, and R. Rigobon. 2010. "Currency Choice and Exchange Rate Pass-Through." *American Economic Review* 100 (1): 304–336.

Hardy, B., F. Saffie, and I. Simonovska. 2024. "Trade credit and exchange rate risk pass through." *BIS Working Paper.*

Jacobson, T., and E. von-Schedvin. 2015. "Trade Credit and the Propagation of Corporate Failure: An Empirical Analysis." *Econometrica.*

Levin, J. 2003. "Relational Incentive Contracts." *American Economic Review.*

Macchiavello, R., and A. Morjaria. 2015. "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports." *American Economic Review.*

Maksimovic, V., and A. Demirgüç-Kunt. 2001. "Firms as Financial Intermediaries: Evidence from Trade Credit Data." *World Bank Working Paper.*

Malcomson, J. 2015. "Relational incentive contracts with productivity shocks." *Games and Economic Behavior.*

Petersen, M., and R. Rajan. 1997. "Trade Credit: Theories and Evidence." *The Review of Financial Studies.*

10 Appendix

Table 14: Supplier-year financial outcomes: non-large suppliers

	dep. variable:		
	ROA_{jt}	Leverage_{jt}	$\text{Profit Margin}_{jt}$
ν_{jt}	-0.068**	0.062**	2.380***
	(0.027)	(0.027)	(0.380)
supplier FE	Yes	Yes	Yes
year FE	Yes	Yes	Yes
obs	233,389	233,389	238,142

Note: Non-large suppliers as defined in Table 2. Standard errors clustered at the supplier level.

Table 15: Robustness: gross exposure measure $\frac{\text{imports of intermediate inputs}}{\text{intermediate inputs}}$

	dep. variable:	
	$\log(\text{quantity}_{ijbt})$	$\log(\text{price}_{ijbt})$
ν_{jt}	-1.664***	0.331
	(0.399)	(0.252)
$\nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t}$	2.045***	-0.416
	(0.428)	(0.318)
seller–buyer–product FE	Yes	Yes
year FE	Yes	Yes
obs	34,519,734	34,519,734

Note: standard errors clustered at the supplier level. This specification uses gross import exposure (not netting out exports) as a robustness check on the preferred net-exposure measure.

Table 16: Robustness: exposure measure $\frac{\text{imports of intermediate inputs}}{\text{sales}}$

	dep. variable:	
	$\log(\text{quantity}_{ijbt})$	$\log(\text{price}_{ijbt})$
ν_{jt}	-1.503***	0.489**
	(0.430)	(0.199)
$\nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t}$	1.441***	-0.025
	(0.315)	(0.170)
seller–buyer–product FE	Yes	Yes
year FE	Yes	Yes
obs	34,519,734	34,519,734

Note: standard errors clustered at the supplier level.

Table 17: Robustness: exposure measure $\frac{\text{imports of intermediate inputs}-\text{exports}}{\text{sales}}$

	dep. variable:	
	$\log(\text{quantity}_{ijbt})$	$\log(\text{price}_{ijbt})$
ν_{jt}	-1.499***	0.493**
	(0.457)	(0.206)
$\nu_{jt} \times \text{TC}^{\text{pre}}_{ij,t}$	1.494***	-0.022
	(0.338)	(0.175)
seller–buyer–product FE	Yes	Yes
year FE	Yes	Yes
obs	34,519,734	34,519,734

Note: standard errors clustered at the supplier level.

Table 18: Excluding 2022

	dep. variable:	
	$\log(\text{quantity}_{ijbt})$	$\log(\text{price}_{ijbt})$
ν_{jt}	-1.673***	0.746***
	(0.782)	(0.197)
$\nu_{jt} \times \text{TC}_{ij,t}^{\text{pre}}$	1.531***	-0.820***
	(0.568)	(0.287)
seller–buyer–product FE	Yes	Yes
year FE	Yes	Yes
obs	24,626,552	24,626,552

Note: standard errors clustered at the supplier level.