Disclosure or Inference?
Information Transmission in Production Networks

Vincenzo Cuciniello, Bank of Italy

Luigi Paciello*, EIEF & CEPR

March 19, 2026

Abstract

We study what governs information transmission in production networks. Combining a randomized information treatment with administrative data on the universe of business-to-business transactions in Italy, we show that firms learn from trading partners through both beliefs and prices, but that the relative importance of these channels varies sharply across relationships.

Belief-based learning is strongest in persistent and economically important supplier relationships, while price-based learning is stronger when suppliers have broader networks. This mirror pattern is driven by supplier-side, rather than buyer-side, relationship characteristics, pointing to supplier disclosure incentives as the key margin governing cross-link heterogeneity in transmission.

We develop a model in which suppliers credibly disclose information about common cost pressures to buyers, eliminating the markup distortion that arises when buyers infer costs from prices. The model predicts that disclosure activates when the buyer–seller information asymmetry is large—when buyers are relatively less informed about suppliers' costs—and that treated suppliers' communicated beliefs—not prices—raise bilateral trade volumes. Both predictions are confirmed in the data. Calibrated to the Italian production network, the model implies that existing information transmission recovers 0.62 percentage points of GDP, predominantly through the markup channel. Aggregate price-level stabilization recovers little of the residual cost: reducing cost volatility causes firms to exit disclosure relationships, crowding out two-thirds of the direct welfare gain.

Keywords: information frictions, production networks, inflation expectations, firm communication, markup distortions
JEL Codes: E31, D83, D85, E52

Corresponding author: luigi.paciello@eief.it. We thank Marios Angeletos, Saroj Bhattarai, Zhen Huo, Yuriy Gorodnichenko, Claudio Michelacci, Jordan Norris and Alireza Tahbaz-Salehi for valuable suggestions and interactions. We benefited from comments received when presenting preliminary versions of this work at EIEF and at the CUHK–CEPR Workshop on "Imperfect Information and Bounded Rationality," Hong Kong.

1 Introduction

Firms in production networks observe local signals—their own costs, prices, and demand—and use them to form beliefs about aggregate states they cannot directly observe, as in Lucas (1972); Lorenzoni (2009); Angeletos and Lian (2023). In such networks, the prices charged by suppliers and the quantities demanded by customers create directed channels through which information can propagate across firms, as studied by La'O and Tahbaz-Salehi (2022). Whether this transmission operates primarily through prices—buyers extracting information statistically—or through beliefs—partners communicating views directly—has distinct implications for markups, quantities, and welfare. Yet direct evidence on learning within production networks remains limited, in part because it is difficult to observe both exogenous variation in firms' information and the bilateral links along which partner-generated signals travel.

This paper studies what governs information transmission in production networks. We combine a unique set of survey, administrative, and experimental data from Italy. Our setting merges the *Survey on Inflation and Growth Expectations* (SIGE), conducted by the Bank of Italy, with administrative records covering the universe of business-to-business transactions between 2019 and 2023. The merged data allow us to observe, for the same firms and quarters, (i) beliefs about aggregate inflation, (ii) realized own price changes, and (iii) the directed trading links along which partner-generated signals may be observed. A central feature of SIGE is a randomized information intervention in which a subset of firms is provided with precise information about recent aggregate inflation before reporting expectations. Because treatment assignment is highly persistent, it generates durable differences in information precision that are orthogonal to firm characteristics and allows trading partners to learn, through repeated interactions, which of their counterparts tend to be better informed about aggregate conditions.

Our first result is that firms learn from better-informed trading partners through both beliefs and prices, but that the relative importance of these channels varies sharply across relationships. Belief-based learning is strongest in durable and economically important supplier relationships. By contrast, price-based learning is stronger when suppliers have broader networks. Thus, direct learning from partners' beliefs and indirect learning from partners' prices move in opposite directions across relationship margins. By contrast, downstream learning from customers operates through the belief channel on average but displays little systematic heterogeneity across relationship margins.

Our second result is that this heterogeneity is fundamentally supplier-side. The gradients are strong and systematic when measured using supplier-side relationship characteristics—relationship persistence, economic importance, and supplier breadth (the number of the supplier's trading partners)—but weaken sharply, and in some cases disappear, when the same characteristics are measured on the buyer side. This asymmetry is difficult to reconcile with buyer-attention explanations, under which own-side relationship characteristics should matter at least as much. It instead points to supplier-originated disclosure as the key margin: whether information is transmitted directly or inferred from prices depends primarily on supplier-side relationship characteristics.

Our third result is that belief-based transmission has real effects. Using administrative transaction data, we show that the belief channel—not the price channel—predicts bilateral trade volumes. When a supplier is better informed, buyers use its beliefs as a signal about cost fundamentals, sustaining bilateral trade. This result shows that the belief channel matters not only for expectations, but also for real buyer–supplier allocations.

To interpret these facts, we develop a model of endogenous information disclosure inside production networks. Suppliers choose whether to communicate information directly to buyers or leave them to infer it from prices. The model combines imperfect information in production networks of La'O and Tahbaz-Salehi (2022) with bilateral communication about private information, in the spirit of Crawford and Sobel (1982). Suppliers can credibly disclose information about common cost pressures to buyers, but truthful communication must be supported by relationship surplus and sufficiently low setup costs. Disclosure occurs in durable, high-surplus relationships and is harder to sustain when suppliers manage broader networks, generating the cross-link heterogeneity documented in the data. When disclosure occurs, buyers no longer need to infer costs from prices, and suppliers charge the frictionless markup. When disclosure fails, buyers rationally extract information from prices, perceived demand elasticity falls, and markups rise above the frictionless level.

This mechanism implies a first-order welfare cost of imperfect communication. That cost is absent from models in which information frictions operate only through rational inattention or noisy signal processing, where welfare effects arise through second-order changes in decision accuracy.[1] The novel margin in our model is not the existence of a markup wedge under price inference, but the endogenous disclosure decision that determines which links rely on direct communication and which remain in the price-inference regime. On the latter links, the markup wedge is a feature of the pricing game rather than of belief formation per se: when buyers infer costs from prices, markups are distorted regardless of how precisely they forecast aggregate inflation. Disclosure shuts down this inference and restores the frictionless markup; no amount of better buyer attention can do this.

The model yields two central empirical predictions. First, disclosure should be more prevalent when buyer–seller informational gaps are more consequential, for example in relationships with higher surplus and lower communication frictions. Second, the belief channel—not the price channel—should be the operative margin for bilateral trade: when a supplier is better informed, buyers can use its communicated beliefs as a credible signal about cost fundamentals, sustaining trade that uncertainty would otherwise depress. We confirm the second prediction using administrative transaction data: a one-percentage-point increase in a treated supplier's inflation expectation raises bilateral trade by approximately 10 percent more than the same increase from an untreated supplier, while the treated-supplier price interaction is small and insignificant conditional on industry controls. Calibrated to the Italian production network, the model implies that existing information transmission recovers 0.62 percentage points of GDP relative to a world without belief channels. Most of this gain operates through the markup channel: disclosure removes the price-inference wedge on non-disclosure links and thereby improves bilateral allocations. Aggregate price-level stabilization delivers only a modest additional gain (0.03 pp), because reducing aggregate cost volatility causes firms to exit disclosure relationships and thereby crowds out much of the direct benefit. The standard second-order channels—aggregate pricing errors and within-industry dispersion—remain operative and would respond to stabilization as in La'O and Tahbaz-Salehi (2022), but they account for a smaller share of the total welfare cost in the calibrated economy.

The model yields two central empirical predictions. First, belief-based transmission should strengthen when the buyer–seller information asymmetry is large, because the value of disclo-

[1] See Mankiw and Reis (2002) and Afrouzi (2024) for leading examples of information-friction models applied to price-setting and inflation dynamics, in which welfare costs arise through second-order changes in decision accuracy.

sure rises with the informational gap between suppliers and buyers—consistent with the state-dependence documented above. Second, the belief channel—not the price channel—should be the operative margin for bilateral trade: when a supplier is better informed, buyers can use its communicated beliefs as a credible signal about cost fundamentals, sustaining trade that uncertainty would otherwise depress. We confirm the second prediction using administrative transaction data: a one-percentage-point increase in a treated supplier's inflation expectation raises bilateral trade by approximately 10 percent more than the same increase from an untreated supplier, while the treated-supplier price interaction is small and insignificant conditional on industry controls. We confirm both predictions in the data. Calibrated to the Italian production network, the model implies that existing information transmission recovers 0.62 percentage points of GDP relative to a world without belief channels, predominantly through the markup channel. This markup cost—arising from buyer–seller information asymmetry about costs, and first order in the size of the asymmetry—responds weakly to aggregate price-level stabilization (0.03 pp gain), because removing aggregate cost volatility causes firms to exit disclosure relationships, crowding out two-thirds of the direct benefit. The standard second-order channels (aggregate pricing errors and within-industry dispersion) remain operative and would respond to stabilization as in La'O and Tahbaz-Salehi (2022), but they account for a smaller share of the total welfare cost in the calibrated economy.

Our paper contributes to a growing literature on how firms form expectations and learn from prices. Andrade et al. (2022) show that firms' aggregate expectations respond to industry-level price variation, consistent with island models. Fang et al. (2025) study firms' endogenous attention allocation within production networks, and Huo et al. (2025) show that heterogeneous attention received by firms—captured by a directed information network—can amplify micro-level shocks into aggregate fluctuations. These papers study how firms allocate attention across broad sets of signals, including competitors and industry leaders, and it may well be that attention to non-transacting firms is an important channel in its own right. Our data, which link survey beliefs to the universe of bilateral transactions, point to a complementary margin: within buyer–seller relationships, the mode of transmission itself is endogenous, and supplier-side communication determines whether information travels through beliefs or through prices.

Our analysis is also related to recent work that uses randomized information provision in surveys to causally identify the effects of expectations on economic decisions. As reviewed in Coibion and Gorodnichenko (2025), this literature primarily exploits one-time or transitory information shocks to trace how agents update from prior to posterior beliefs and how these belief revisions affect subsequent choices.[2] The closest empirical paper to ours is Hajdini et al. (2025), who use a randomized information treatment in New Zealand to show that beliefs and actions propagate across buyer–supplier links, that the belief channel is stronger in durable and economically important relationships, and that these gradients are driven by supplier-side rather than buyer-side relationship characteristics. Our setting allows us to go beyond propagation in three dimensions. First, because the Italian treatment is highly persistent, we can identify which links are durably in the disclosure regime rather than capturing one-time belief revisions. Second, linking the survey to the universe of bilateral transactions allows us to test whether information transmission distorts the volume of trade between specific buyer–supplier pairs—the margin that generates the first-order welfare cost in our model and that is distinct from whether firms adjust their own prices or employment in response to partners' actions. Third, we develop a model

[2]Weber et al. (2025) provide cross-country evidence from survey experiments, showing that learning from information treatments differs markedly between low- and high-inflation environments.

with an explicit incentive constraint that delivers the disclosure threshold, quantify the welfare cost of price inference, and show that aggregate stabilization cannot eliminate it.

On the theory side, our framework combines three strands of work. The first studies learning from prices in networks: Babus and Kondor (2018) map decentralized price formation into a conditional guessing game, La'O and Tahbaz-Salehi (2022) study price setting and monetary policy in production networks under imperfect information, Nikolakoudis (2024) analyze learning from prices that emerge through market clearing, and Bui et al. (2024) study how noisy information propagates through global value chains, generating higher-order volatility along supply-chain layers. The second studies the informational role of prices in vertical relationships: Gaballo and Paciello (2025) show that when buyers infer costs from supplier prices, the supplier's perceived demand elasticity falls and markups are distorted; Janssen and Shelegia (2015) demonstrate that unobserved wholesale prices create inference problems that worsen double marginalization. The third studies how the public or private nature of information shapes equilibrium outcomes—Morris and Shin (2002) show that publicly observed signals can crowd out beneficial private information; Amador and Weill (2010) show in a macro model with price-based learning that public information releases can reduce the informational efficiency of the price system; and Amador and Weill (2012) show that when agents learn through both a public channel (prices) and a private channel (bilateral interactions), more public information can reduce asymptotic knowledge—and how communication can be sustained inside bilateral relationships when information transmission is not automatic. Our contribution is to combine these insights in a framework where the mode of information transmission itself is endogenous: supplier-side communication determines which links transmit information through beliefs and which remain in the price-inference regime.

More broadly, our paper relates to work on information frictions, coordination, and network propagation in macroeconomics, including Acemoglu et al. (2012), Acemoglu and Tahbaz-Salehi (2025), Gabaix (2014), Angeletos and Huo (2021), and Ghassibe and Nakov (2025); see Carvalho and Tahbaz-Salehi (2019) for a survey. Relative to this broader literature, our contribution is to identify a concrete microeconomic margin—whether suppliers communicate directly or leave buyers to infer from prices—and to quantify its implications for real allocations and welfare inside production networks.

The remainder of the paper is organized as follows. Section 2 describes the survey data on expectations and prices, the administrative transaction records, and the construction of the firm-level production network. Section 3 presents the first-stage evidence on treatment and forecast accuracy (Section 3.1) and the average learning regressions (Section 3.2). Section 4 studies heterogeneity across relationship size, persistence, and partner breadth, documents the contrast between partner-side and own-side gradients, and reports an additional state-dependent pattern in information transmission. Section 5 develops the model of strategic information disclosure in production networks. Section 6 takes the model to the data: it calibrates the model to the experimental moments, characterizes the equilibrium disclosure set, evaluates the model's bilateral trade implications, and reports welfare decompositions across counterfactual information regimes. Section 7 concludes.

2 Data

Our empirical analysis combines firm-level survey data on inflation expectations and firm price growth with administrative data on business-to-business transactions. The merged dataset al-

lows us to observe, for the same firms and quarters, (i) beliefs about aggregate inflation, (ii) realized own price changes, (iii) directed buyer–supplier relationships, and (iv) a rich set of firm and partner characteristics. For each surveyed firm, the data include information on sector, geographic location, employment size, and other time-invariant and time-varying attributes reported in the survey.[3] For each linked trading partner, the administrative records provide identifiers, sectoral classification, and transaction values, which allow us to characterize the economic importance of each link. This combined dataset yields a quarterly panel in which surveyed firms are embedded in observed production networks, linking firms' beliefs and prices to those of their trading partners. This section proceeds as follows. Section 2.1 describes the survey data on expectations and prices and the randomized information treatment. Section 2.2 details the administrative transaction data and the construction of buyer–supplier networks.

2.1 Survey data on expectations and prices

The primary source of expectations data is the quarterly *Survey on Inflation and Growth Expectations* (SIGE) conducted by the Bank of Italy. SIGE is a quarterly stratified survey of Italian firms with at least 50 employees operating in industry, nonfinancial private services, and construction, with sampling designed to ensure coverage across sector, firm size, and geographic region. Each wave includes roughly 1,000–1,300 firms, and more than 5,900 firms have participated since 1999. The survey elicits quantitative expectations of Italian CPI inflation at multiple horizons (6 months, 12 months, and, in later waves, longer horizons), as well as the firm's average year-on-year own price change. In addition, SIGE records firm characteristics—including employment, sectoral classification, and geographic location—which we use to condition on observable heterogeneity.[4] A central feature of our identification strategy is a randomized information treatment embedded in SIGE since 2012. In each wave, a randomly selected subset of firms—approximately two-thirds of participants—is provided with publicly available information on recent CPI inflation in Italy and the euro area immediately before reporting their inflation expectations; the remaining firms receive no such information.[5] Assignment to treatment is persistent within re-randomization windows, but not permanent over the entire life of the survey. Treatment status was first randomized in 2012:Q4 and remained fixed until 2017:Q2, when assignment was redrawn Coibion et al. (2020). In 2019:Q4, approximately 14 percent of firms were reassigned from treatment to control; treatment status is otherwise unchanged through 2023:Q4.[6] Our main linked buyer–supplier sample covers 2019–2023, a period during which treatment status is predetermined and time-invariant for the vast majority of firms. This persistence allows trading partners to learn, through repeated interactions, which counterparts are systematically better informed about aggregate conditions.

Figure 1 illustrates the evolution of inflation expectations in the SIGE survey for treated and

[3] A detailed description of the SIGE survey is provided in the supplementary material to Coibion et al. (2020). Rosolia (2024) uses the same survey to study whether Italian firms act on their inflation expectations.

[4] Appendix J provides additional details on the SIGE survey design, sampling, variable construction, and the information treatment protocol. The full questionnaire is reproduced in the Online Appendix to Coibion et al. (2020). The SIGE microdata are made available by the Bank of Italy to external researchers through a controlled-access procedure; information on data access and application requirements is available at Bank of Italy BIRD microdata access.

[5] Coibion et al. (2020), Online Appendix Table 2, show that treatment assignment is uncorrelated with firm size, sector, geographic area, export intensity, and pricing behavior.

[6] Results are unchanged when we exclude reassigned firms and their trading partners (Appendix U).

untreated firms, together with realized CPI inflation. The figure plots average expectations of CPI inflation two quarters ahead against the corresponding realized inflation rate. Treated firms (dashed line) exhibit expectations that track realized inflation (dotted line) more closely than untreated firms (solid line) throughout the sample period. The vertical distance between each expectations series and realized inflation provides a direct visualization of forecast errors. The persistent gap between treated and untreated firms' forecast accuracy motivates our analysis of whether this information advantage propagates through production networks.

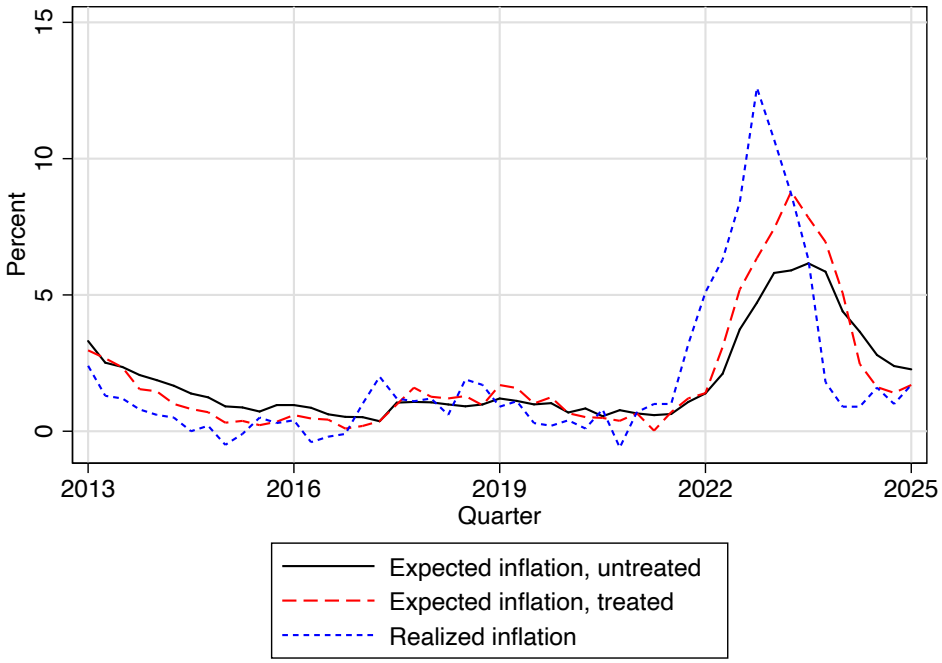


Figure 1: Inflation expectations two quarters ahead and realized inflation

Notes: The figure plots average expected CPI inflation two quarters ahead as reported in the SIGE survey by untreated firms (solid line) and treated firms (dashed line), together with realized CPI inflation (dotted line) in the corresponding quarter. Expectations are aggregated by quarter. Treated firms receive information on recent CPI inflation prior to reporting expectations.

2.2 Administrative transaction data and network construction

We link SIGE firms to administrative data covering the universe of business-to-business transactions in Italy from 2019 to 2023. These records originate from mandatory electronic invoicing, whose scope was expanded on 1 January 2019 to include all transactions in goods and services between VAT-registered entities resident in Italy.[7] The underlying records report transaction-level values and dates. We aggregate transactions to the quarterly frequency to align them with the timing of SIGE.

Using these data, we construct quarterly buyer–supplier networks. For each SIGE firm i and quarter t, we observe upstream links to suppliers and downstream links to customers, together

[7]Appendix K provides additional details on the electronic invoicing data, the matching procedure, and the construction of network variables. The electronic invoicing data are provided by the Italian Revenue Agency (*Agenzia delle Entrate*) and accessed through the Bank of Italy's internal data infrastructure.

with the value of transactions aggregated over the quarter along each link.

Our baseline empirical unit is a quadruple (i, j, t, d), where i and j are firms linked by at least one transaction in quarter t, and $d \in \{\text{up}, \text{dn}\}$ denotes the direction of the relationship, indicating whether j is a supplier of firm i (upstream) or a customer of firm i (downstream). The direction distinction is important because the outcome-based signal differs across the two sides of a trading relationship: a buyer can observe its supplier's posted price, whereas a seller typically observes quantities or revenues rather than the customer's own price. We therefore define direction-specific outcome signals—supplier price growth upstream, customer purchase growth downstream—and allow all spillover coefficients to vary by direction throughout the empirical analysis.

The analysis sample is restricted to trading relationships in which both firms are observed in SIGE in the same quarter, since expectations and price variables are available only for surveyed firms. For the cross-sectional analysis in Section 3, we focus on fourth-quarter observations to provide a consistent annual snapshot and to mitigate seasonality concerns.

Firms in the SIGE sample are economically sizable and embedded in extensive production networks. The median firm employs roughly 100 workers and trades with approximately 191 suppliers and 77 buyers in the universe of administrative transactions. However, only a small fraction of these trading partners also participate in SIGE: at the median firm–quarter, surveyed suppliers account for about 2.6% of all upstream relationships, typically corresponding to 3–5 observed suppliers per firm–quarter. Despite this sparse coverage, most firms are matched to multiple surveyed partners, generating meaningful cross-sectional variation in network exposure.

The estimation dataset is constructed in a directional format: each matched firm–partner–quarter observation appears twice, once for each trade direction, yielding a balanced sample across directions. In total, this corresponds to 179,464 distinct matched trading links, or 358,928 directional observations.

Table 1 summarizes the structure of matched trading links. In each fourth quarter, between 1,020 and 1,251 firms have at least one matched upstream supplier, compared to 780–1,004 with at least one matched downstream customer. The asymmetry reflects differences in firms' network structures: although each trading link appears symmetrically in the data, firms differ in how many suppliers versus customers they have and in which partners overlap with SIGE.

Conditional on having at least one match, firms are typically connected to several surveyed partners: the median firm–quarter is matched to 6 upstream suppliers and 4 downstream customers, with firms in the upper quartile linked to 11 upstream and 9 downstream partners. The last columns of Table 1 report the share of total quarterly trade value accounted for by individual matched partners. At the median, matched upstream links cover about 1–1.5% of total input purchases, rising to roughly 7–11% at the 75th percentile, with similar magnitudes downstream. These coverage rates reflect sparse but economically meaningful sampling of firms' production networks. A limitation of this design is that both sides of every matched link are SIGE participants—firms with at least 50 employees. The information environment in relationships between large surveyed firms may differ from that in relationships with smaller, unsurveyed partners, where communication infrastructure is less developed and managerial attention more constrained. Our estimates therefore characterize learning within the large-firm production network; whether the same patterns extend to smaller firms is an open question.

Table 1: Matched trading partners per firm, by quarter and direction

	# of matches							% of value	
Quarter	≥ 1	QI	QII	QIII	QI*	QII*	QIII*	QI	QII	QIII
Upstream										
2019:Q4	1,020	3.00	5.00	9.00	1.00	3.00	6.00	0.2	1.5	11.1
2020:Q4	1,223	3.00	6.00	10.00	2.00	3.00	6.00	0.2	1.4	10.0
2021:Q4	1,236	4.00	7.00	12.00	2.00	4.00	8.00	0.1	1.0	7.4
2022:Q4	1,251	4.00	7.00	11.00	2.00	4.00	8.00	0.2	1.2	7.9
2023:Q4	1,242	4.00	6.00	11.00	2.00	5.00	8.00	0.1	1.1	7.8
Downstream										
2019:Q4	780	2.00	3.50	8.00	1.00	2.00	5.00	0.1	1.1	7.8
2020:Q4	910	2.00	3.00	8.00	1.00	2.00	5.00	0.1	1.1	7.8
2021:Q4	961	2.00	4.00	9.00	1.00	2.00	6.00	0.1	1.1	7.8
2022:Q4	1,004	2.00	4.00	10.00	1.00	2.00	7.00	0.1	1.1	7.8
2023:Q4	985	2.00	4.00	10.00	1.00	3.00	7.00	0.1	1.1	7.8

Notes: The table is computed using Q4 observations only. Column ≥ 1 reports the number of SIGE firms with at least one matched trading partner in the given direction and quarter. QI, QII, and QIII denote the 25th, 50th, and 75th percentiles of the number of matched partners per firm. Columns marked with an asterisk restrict attention to treated partners. The last three columns report the 25th, 50th, and 75th percentiles of the share of a firm's total quarterly trade value accounted for by an individual matched partner in the same direction. Shares are computed within firm, quarter, and direction and are expressed in percentage points.

3 Learning from trading partners

This section presents the main empirical evidence on information transmission across buyer-supplier links. We proceed in three steps. Section 3.1 establishes that the information treatment improves firms' inflation forecasts and is reflected in their realized prices, making treated firms more informative trading partners. Section 3.2 then shows that firms learn from better-informed partners through both beliefs and prices. Finally, Section 4 shows that the balance between these channels varies systematically across relationships, and that the strongest heterogeneity is driven by supplier-side rather than buyer-side characteristics. This supplier-side asymmetry is the central empirical fact motivating the disclosure mechanism developed in Section 5.

3.1 First-stage evidence: treatment and forecast accuracy

This subsection establishes that the information treatment improves firms' inflation forecasts, providing the first-stage foundation for the network learning analysis that follows.[8] The result motivates both channels of information transmission studied in Section 3.2. Because treated firms hold more accurate beliefs, their stated inflation expectations are more informative signals for trading partners—the belief channel. Because these more accurate beliefs are reflected in the prices treated firms set, their realized price changes also carry more precise information about aggregate inflation—the price channel.

[8]The survey asks firms to report the *"Variazione percentuale media annua dei prezzi dei beni da Voi fatturati"*, which translates as the average year-on-year percentage change in the prices of the goods billed by the firm.

The analysis proceeds in two steps. First, we show that treated firms' survey expectations are closer to realized inflation. Second, we construct a price-based diagnostic that isolates the component of expectations predictable from own price growth and show that this component is also more accurate for treated firms. Together, these results imply that the treatment-induced information advantage is embedded in both firms' stated beliefs and their pricing decisions, making treated firms more informative partners through both the belief and price channels. Because treatment assignment is persistent and the panel spans both low- and high-inflation regimes (2013–2023), we can also characterize the time profile of the first-stage effect and show that the informational advantage of treated firms rises sharply when aggregate inflation increases—a pattern that motivates the State-dependent transmission mechanism developed in Section 5.

Let π_t denote realized year-on-year Italian CPI inflation in quarter t, and let $\pi_{i,t}^{e,h}$ denote firm i's h-quarters-ahead inflation expectation recorded in SIGE in quarter $t-h$. By construction, $\pi_{i,t}^{e,h}$ is a forecast of π_t. Let $T_i \in \{0,1\}$ indicate whether firm i is assigned to the information treatment. Throughout this subsection, we work with winsorized variables (1st/99th percentiles, weighted) and use the full SIGE panel 2013:Q1–2023:Q4.[9] Our baseline measure of forecast accuracy is the absolute forecast error

$$AFE_{i,t}^{\pi^e} \equiv \left| \pi_t - \pi_{i,t}^{e,h} \right|, \tag{1}$$

indexed by realization quarter t and evaluating expectations formed in quarter $t-h$.

To assess whether the treatment also improves the informativeness of firms' pricing decisions, we construct a price-based diagnostic. Let $g_{i,t-h}^{p}$ denote firm i's realized year-on-year own price growth reported in SIGE in quarter $t-h$. For each quarter t and treatment group $T \in \{0,1\}$, we estimate the cross-sectional projection

$$\pi_{i,t}^{e,h} = a_{t,T} + b_{t,T}\, g_{i,t-h}^{p} + \varepsilon_{i,t}, \tag{2}$$

and define the price-based forecast error as

$$AFE_{i,t}^{\hat{\pi}^e} \equiv \left| \pi_t - \hat{\pi}_{i,t}^{e,h} \right|, \tag{3}$$

where $\hat{\pi}_{i,t}^{e,h} \equiv a_{t,T} + b_{t,T}\, g_{i,t-h}^{p}$. This measure isolates the component of inflation expectations predictable from own price growth. If the treatment improves $AFE_{i,t}^{\hat{\pi}^e}$, it follows that treated firms' prices embed more precise information about aggregate inflation at the time expectations are formed.

We estimate

$$Y_{i,t} = \alpha_t + \alpha_i + \beta\, T_{i,t} + \varepsilon_{i,t}, \tag{4}$$

where α_t are quarter fixed effects, α_i (included in some specifications) are firm fixed effects, and $T_{i,t}$ is firm i's treatment indicator in quarter t. All regressions use SIGE sampling weights and cluster standard errors at the firm level.

We consider three dependent variables. The first is the survey-based absolute forecast error $AFE_{i,t}^{\pi^e}$ defined in (1). The second is the in-sample price-based diagnostic $AFE_{i,t}^{\hat{\pi}^e}$ defined in (3), which projects expectations on contemporaneous own price growth within each quarter and treatment group. The third is a rolling-window variant, $AFE_{i,t}^{\hat{\pi}^e,\text{roll}}$, which estimates the same projection using only data available prior to quarter $t-h$, ruling out mechanical in-sample fit.

Table 2 reports estimates of β for 2-quarters-ahead expectations ($h=2$).[10] Columns 1–3

[9] The network learning regressions in Sections 3.2 and 4 are restricted to 2019–2023, when the administrative transaction data become available. We use the longer window here to maximize precision and to characterize the treatment effect across both low- and high-inflation regimes.

[10] Similar results obtain at alternative horizons.

include quarter fixed effects; columns 4–6 add firm fixed effects.

Across all specifications, the treatment significantly reduces absolute forecast errors. Treated firms have survey-based forecast errors 0.58 percentage points lower than control firms (column 1), and 0.48 pp lower with firm fixed effects (column 4). The same pattern holds for both price-based diagnostics. Under the rolling-window specification, which uses only past data, the treatment effect ranges from 0.27 to 0.34 pp (columns 3 and 6). Treated firms' prices therefore embed more precise information about aggregate inflation at the time expectations are formed, establishing the first-stage foundation for both the belief and price channels studied in the network regressions that follow.

Table 2: Treatment and absolute forecast errors

	Time FE			Firm + Time FE		
	$AFE^{\pi^e}_{i,t}$	$AFE^{\hat{\pi}^e}_{i,t}$	$AFE^{\hat{\pi}^e,\text{roll}}_{i,t}$	$AFE^{\pi^e}_{i,t}$	$AFE^{\hat{\pi}^e}_{i,t}$	$AFE^{\hat{\pi}^e,\text{roll}}_{i,t}$
T_i	-0.581^{***}	-0.543^{***}	-0.340^{***}	-0.483^{***}	-0.453^{***}	-0.270^{***}
	(0.026)	(0.006)	(0.006)	(0.039)	(0.007)	(0.008)
N	46,936	46,936	45,906	46,331	46,331	45,297

Notes: Dependent variable is absolute forecast error (AFE). All regressions are weighted using analysis weights (survey weight if available, otherwise 1). Forecast-error measures are defined relative to realized inflation in quarter t and expectations formed in quarter $t-h$. Columns 1–3 include quarter fixed effects; columns 4–6 add firm fixed effects, so that β is identified from within-firm changes in treatment status at the 2017:Q2 and 2019:Q4 re-randomization dates. AFE^{π^e} uses survey expectations directly. $AFE^{\hat{\pi}^e}$ uses the component predictable from own price growth via in-sample cross-sectional projection within treatment group and quarter. $AFE^{\hat{\pi}^e,\text{roll}}$ estimates the projection using only past data within each treatment group. Expectations are 2 quarters ahead. All variables are winsorized at the 1st and 99th percentiles (weighted). Standard errors are clustered at the firm level.

To characterize the time profile of the treatment effect, we estimate quarter-specific cross-sectional regressions

$$Y_{i,t} = \alpha + \beta_t T_{i,t} + u_{i,t}, \qquad Y_{i,t} \in \left\{ AFE^{\pi^e}_{i,t},\, AFE^{\hat{\pi}^e}_{i,t} \right\}, \tag{5}$$

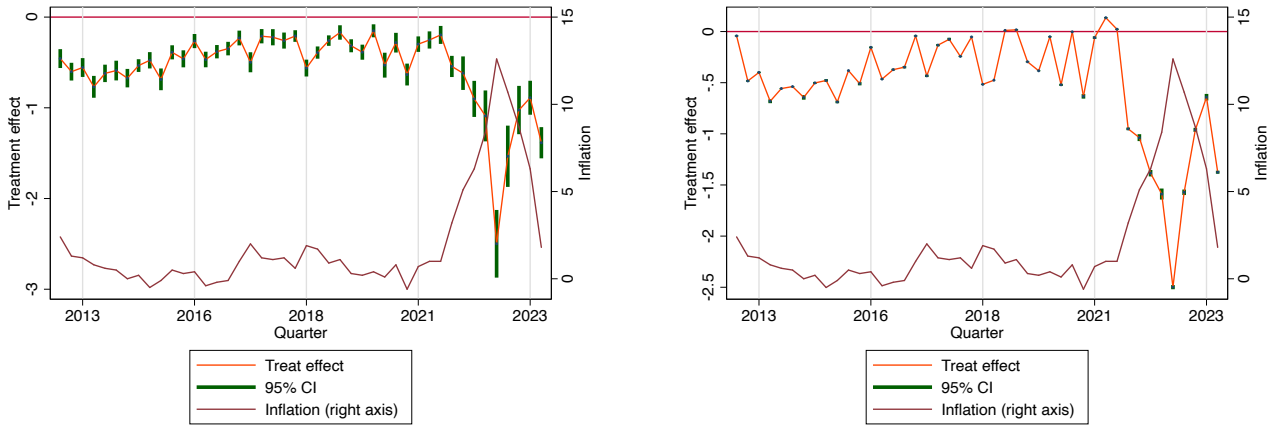
separately for each quarter t. Figure 2 plots the estimated $\hat{\beta}_t$ together with 95% confidence intervals and realized Italian CPI inflation.

Two patterns stand out. First, the treatment effect is negative in essentially every quarter for both outcomes, confirming that the average effects in Table 2 are not driven by a subset of periods. Second, the magnitude of the treatment effect increases sharply when aggregate inflation rises after 2021. In both panels, $|\hat{\beta}_t|$ tracks the acceleration in CPI inflation closely. The amplification is especially pronounced for the price-based diagnostic, whose mechanical construction filters out firm-specific expectation noise and yields smoother, more precisely estimated coefficients.[11]

[11] Standard errors on the price-based diagnostics are an order of magnitude smaller than those on $AFE^{\pi^e}_{i,t}$ because the projection in (2) absorbs idiosyncratic variation in expectations, leaving a dependent variable with substantially less cross-sectional dispersion.

Figure 2: Quarter-specific treatment effects and inflation

(a) $Y_{i,t} = AFE^{\pi^e}_{i,t}$ **(b)** $Y_{i,t} = AFE^{\hat{\pi}^e}_{i,t}$



Notes: Each panel plots $\hat{\beta}_t$ from equation (5), with 95% confidence intervals (shaded). Dashed line: realized Italian CPI inflation (right axis). Left panel: survey-based forecast errors ($AFE^{\pi^e}_{i,t}$). Right panel: price-based diagnostic ($AFE^{\hat{\pi}^e}_{i,t}$). Regressions use SIGE sampling weights; standard errors clustered at the firm level.

3.2 Learning through beliefs and prices

We now test whether firms learn from better-informed trading partners, and through which channel. We distinguish between two modes of transmission. The first is a belief channel, in which firms update their inflation expectations in response to partners' stated beliefs. The second is a price channel, in which firms extract information from partners' observed prices or demand conditions. Our goal is not only to establish that both channels operate, but also to compare their relative importance across buyer-supplier links.

The first-stage results establish that treated firms are better informed and that this advantage is reflected in their prices. Consistent with this, treatment also reduces expectation disagreement between linked firms (Appendix B). The key result of this subsection is that spillovers operate primarily through partner beliefs. In both upstream and downstream relationships, firms place significantly greater weight on the stated inflation expectations of treated partners than on those of untreated partners. By contrast, learning from transaction-based signals is much weaker on average. These average effects are informative in their own right, but they also set up the sharper fact that motivates the rest of the paper: the balance between belief-based transmission and price-based inference varies systematically across relationships, and that heterogeneity is overwhelmingly supplier-side.

The unit of observation is a dyad–quarter–horizon cell constructed from the matched survey–transaction data. For each link, we observe two partner-generated signals. The first is partner j's stated inflation expectation, $\pi^{e,h}_{j,t}$, which captures the belief channel. The second is a transaction-based signal, $g_{j,t}$, whose definition depends on the direction of the relationship. Upstream, the natural signal is supplier j's year-on-year own price growth $g^p_{j,t}$, since buyers observe the prices charged by suppliers. Downstream, the natural signal is customer j's year-on-year purchase growth $g^r_{j,t}$, since sellers observe the demand conditions revealed by customers' purchases.[12]

[12]Purchase growth $g^r_{j,t}$ is the year-on-year growth rate of the total taxable amount (*imponibile*) invoiced to

A maintained assumption is that firms can assess which trading partners tend to be more informative. This does not require knowledge of the SIGE treatment protocol. It requires only that, through repeated interaction, firms learn which partners provide more accurate signals about aggregate conditions. This is plausible in our setting because treatment assignment is highly persistent, so relative forecast quality is stable over time, and because the first-stage results imply a substantial treatment-induced gap in forecast accuracy. Any residual misclassification attenuates the estimated differential response to treated partners toward zero.

To increase precision, we pool the $h = 6$ and $h = 12$ expectation horizons by stacking the two observations at the firm–partner–quarter level and imposing common slope coefficients across horizons.[13] We include a lag of firm i's horizon-matched expectation, $\pi^{e,h}_{i,t-1}$, to absorb persistence in beliefs.

The estimation is performed jointly across upstream and downstream relationships in a single regression. For each dyad (i, j, t) with $j \in \mathcal{N}_{i,t}$, the transaction-based regressor is direction-specific:

$$g_{j,t} = \begin{cases} g^p_{j,t} & \text{if } j \text{ is a supplier of } i \text{ in quarter } t, \\ g^r_{j,t} & \text{if } j \text{ is a customer of } i \text{ in quarter } t. \end{cases}$$

The asymmetry in signal definitions reflects the different information that transactions reveal to each side of a trading relationship. Upstream, a buyer observes the prices charged by its suppliers; the relevant transaction-based signal is therefore the supplier's price growth $g^p_{j,t}$. Downstream, a seller observes the quantities demanded by its customers; the relevant signal is customer purchase growth $g^r_{j,t}$, which captures shifts in the demand conditions faced by the partner.

Ideally, these signals would be measured at the bilateral level—the price that supplier j charges specifically to buyer i, or the purchases that customer j directs specifically at seller i—since these are the magnitudes firms observe in their own transactions. While individual transactions between i and j are recorded in the administrative data, they do not occur in every period, which prevents us from constructing bilateral year-on-year growth rates. We therefore use partner j's aggregate price growth $g^p_{j,t}$ and aggregate purchase growth $g^r_{j,t}$ as proxies, introducing classical measurement error that attenuates estimated coefficients on the transaction-based signal toward zero. This should be kept in mind when interpreting the relative magnitudes of the price and belief channels. As a robustness check, we also estimate a symmetric specification that uses supplier price growth $g^p_{j,t}$ for both directions, so that the transaction-based signal is the partner's own price growth regardless of whether the partner is upstream or downstream. The belief-channel results are essentially unchanged (Appendix L).

All continuous regressors are winsorized at the 1st and 99th percentiles. The partner set $\mathcal{N}_{i,t}$ is defined by contemporaneous transactions: firm j belongs to $\mathcal{N}_{i,t}$ if and only if a transaction between i and j is recorded in the administrative data during quarter t. This ensures that the signals we attribute to firm i's information set are generated by partners with whom i has an active commercial relationship—and hence direct exposure—in that period.

Regressions use analytic weights that allocate each firm–quarter SIGE sampling weight across observed partners. Let $w_{i,t}$ denote the sampling weight. Each dyad (i, j, t) receives weight $\omega_{ij,t} = w_{i,t}\, s_{ij,t}$, where $s_{ij,t}$ is a within-firm link share satisfying $\sum_{j \in \mathcal{N}_{i,t}} s_{ij,t} = 1$, so that the total weight attached to firm i in quarter t equals $w_{i,t}$ regardless of the number of observed links.

Our baseline uses uniform link shares (*unif*; $s_{ij,t} = 1/|\mathcal{N}_{i,t}|$), which treat all contemporaneous

customer j, computed from the administrative transaction data.

[13] Appendix L reports horizon-specific estimates.

partners equally. We report robustness to three alternatives. Breadth weights (*brd*; $s_{ij} \propto n_j^{\text{dir}}$, the number of j's direction-specific trading partners—j's suppliers for upstream links, j's customers for downstream links—summed over the full sample) upweight signals from well-connected partners, who aggregate information from a larger set of counterparts. Transaction-value weights (*val*; $s_{ij} \propto \bar{v}_{ij}$, the bilateral taxable amount cumulated over the full sample) upweight economically larger relationships. Relationship-intensity weights (*shr*; $s_{ij,t} \propto v_{ij,t}/V_{j,t}$, the contemporaneous bilateral taxable amount divided by partner j's total sales in quarter t) give more weight to links in which firm i accounts for a larger share of j's current revenue. Breadth and transaction-value link scores are computed once over the full sample and are therefore time-invariant; uniform and relationship-intensity scores are quarter-specific. In all four cases, shares are renormalized within each firm–quarter–direction cell to sum to one.

The pooled-horizon, joint-direction estimating equation is estimated on the set of all dyads (i, j, t) with $j \in \mathcal{N}_{i,t}$:

$$\begin{aligned}\pi_{i,t}^{e,h} = \alpha_i + \alpha_h + \alpha_{s(i)\times t} + \alpha_{d(i)\times t} + \alpha_{a(i)\times t} + \alpha_{c(i,j)\times t} + \rho\, \pi_{i,t-1}^{e,h} \\ + f(\, g_{j,t},\ T_{i,t},\ T_{j,t}\,;\, \boldsymbol{\beta}\,) + f(\, \pi_{j,t}^{e,h},\ T_{i,t},\ T_{j,t}\,;\, \boldsymbol{\delta}\,) + \varepsilon_{ij,t},\end{aligned} \tag{6}$$

where $f(x, T_{i,t}, T_{j,t}; \boldsymbol{\theta}) = \theta_0\, x + \theta_1\, T_{j,t}\, x + \theta_2\, T_{i,t}\, x + \theta_3\, T_{i,t}\, T_{j,t}\, x$ is a full interaction of signal x with own and partner treatment status. Here α_i is a firm fixed effect, α_h a horizon fixed effect, and $\alpha_{s(i)\times t}$, $\alpha_{d(i)\times t}$, $\alpha_{a(i)\times t}$ are sector×quarter, size×quarter, and area×quarter fixed effects—the same interacted controls used in Coibion et al. (2020) to absorb group-level time variation in the SIGE sample. The term $\alpha_{c(i,j)\times t}$ denotes direction×treatment-pair×quarter fixed effects, included in the full-sample specifications only, which absorb the common component of partner signals within each direction–treatment cell. The coefficients $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$ are identified from idiosyncratic variation in partner signals—variation that reflects either noise in partners' beliefs or partner-specific shocks to prices and demand. A positive loading on this variation indicates that firms extract information from the partner-specific component of observed signals, in the spirit of the signal extraction problem studied by Lucas (1972). Appendix A derives this interaction structure as a first-order approximation to the model's comparative static: when partner j's information precision improves (treatment), the projection coefficients in firm i's posterior shift, generating the $T_{j,t} \times x_{j,t}$ interactions in the regression. The direction×treatment-pair×quarter fixed effects play a key role in this interpretation. They absorb the direct effect of own treatment on expectations—the object studied by Coibion et al. (2020)—as well as the average effect of partner treatment status, which in existing network designs is estimated jointly with the direct effect Hajdini et al. (2025). Our coefficients identify a further margin: whether firms respond differentially to the *realized idiosyncratic signals* generated by treated versus untreated partners, after removing the common component shared by all partners in the same treatment-direction-quarter cell.

Table 3 reports the pooled-horizon estimates from the joint regression in equation (6), with upstream (supplier) and downstream (customer) signals organized in Panels A and B, respectively. Columns (1)–(4) report specifications estimated on untreated firms only ($T_{i,t} = 0$), under the baseline uniform weights and three alternatives (breadth-weighted, transaction-value weighted, and partner-sales-share tilted). Columns (5)–(8) report the full-sample specifications under the same sequence of weight schemes, allowing treated firms to load differently on treated partners' signals. Since both panels are extracted from a single joint regression, the persistence parameter ρ and fit statistics are common across directions. Standard errors use the Driscoll

and Kraay (1998) correction with bandwidth 2 throughout.[14]

Belief spillovers dominate on average. Across both directions, the most robust evidence of information transmission operates through partner beliefs. On customer links, δ_1 ($T_{j,t}\,\pi^{e,h}_{j,t}$) is positive and precisely estimated across all weighting schemes in the untreated-only specifications (0.064 to 0.113, Panel B, columns 1–4) and remains positive in the full-sample specifications (0.050 to 0.098, columns 5–8). On supplier links, δ_1 is also positive and statistically significant under uniform, transaction-value, and tilted weights in both blocks (0.026 to 0.060 in columns 1–4 and 0.032 to 0.069 in columns 5–8, Panel A). In both directions, the level effect of untreated-partner beliefs (δ_0) is small and imprecise. On average, firms therefore update primarily in response to the stated beliefs of better-informed partners.

Price-based inference is weaker and noisier. Learning from transaction-based signals is much weaker on average. Upstream, β_1 ($T_{j,t}\,g_{j,t}$) is positive but imprecisely estimated under most weighting schemes, reaching significance only under breadth weights in the full-sample specification (0.013, Panel A, column 6). Downstream, coefficients on customer purchase growth and its interaction with treatment are economically small, with significance appearing only under the tilted-weight untreated-only specification (0.002, Panel B, column 4) and no stable pattern in the full-sample columns. This asymmetry is consistent with direct communication being a cleaner and less noisy vehicle for transmission than inference from prices or demand conditions.

Own information reduces reliance on partner signals. In the full-sample specifications, treated firms rely less on treated partners' belief signals. The coefficient δ_3 ($T_{i,t}\,T_{j,t}\,\pi^{e,h}_{j,t}$) is negative for both directions and statistically significant under uniform weights upstream and downstream (-0.034 and -0.066, respectively; Panels A and B, column 5). For the transaction-based signal, the analogous term β_3 is negative and significant upstream under breadth weights (-0.015; Panel A, column 6) and is otherwise small. Taken together, these patterns are consistent with own and partner information being substitutes: firms place less weight on partner signals, especially partner beliefs, when their own information is more precise.

The belief-channel coefficients are robust across weighting schemes, whereas the coefficients on the transaction-based signal are less precisely estimated and become most visible under breadth weights. This pattern is consistent with price-based inference being a weaker and noisier signal than direct communication. Appendix T shows that the results are robust to richer partner controls and alternative fixed-effect structures.

Appendix I reports two complementary assessments: a variance decomposition (partner beliefs add 0.18 percentage points to R^2, concentrated in the customer direction) and a linear counterfactual treating all untreated partners (mean expectation shift of 0.041 percentage points, roughly 4% of the typical quarterly update). These reduced-form magnitudes are modest, as expected for experimental variation that perturbs precision for a subset of partners; Section 6 embeds this mechanism in the full production network.

These average effects conceal the central empirical fact of the paper: the balance between

[14]The Driscoll–Kraay estimator accounts for both cross-sectional dependence and serial correlation in the residuals. The bandwidth of two quarters accommodates the overlap between the six- and twelve-month forecast horizons in the stacked panel. In specifications where the Driscoll–Kraay variance matrix is not positive definite—typically because absorbing high-dimensional fixed effects leaves too few effective time-series observations—we fall back to clustering at the firm level, which is strictly more conservative for cross-sectional dependence. In practice, this fallback binds only for the most saturated fixed-effect structures (e.g., the tilted-weight specifications with direction×treatment-pair×quarter effects). All main results are robust to using firm-level clustering throughout. Appendix L reports the same column structure estimated separately at $h = 6$ and $h = 12$, and under quarterly-normalized link-weight normalizations.

Table 3: Learning from supply-chain partners

	Untreated only ($T_{i,t}=0$)				Full sample			
	(1) unif	(2) brd	(3) val	(4) shr	(5) unif	(6) brd	(7) val	(8) shr
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^{p}_{j,t}$)								
β_0: $g_{j,t}$	-0.002	-0.004	-0.003	0.006	-0.004	-0.006	-0.005	0.006
	(0.004)	(0.004)	(0.005)	(0.009)	(0.005)	(0.004)	(0.006)	(0.009)
β_1: $T_{j,t}\, g_{j,t}$	0.010	0.009	0.001	-0.002	0.012	0.013*	0.004	-0.002
	(0.007)	(0.006)	(0.005)	(0.007)	(0.008)	(0.006)	(0.006)	(0.007)
β_2: $T_{i,t}\, g_{j,t}$					0.005	0.006	0.008	-0.007
					(0.006)	(0.005)	(0.006)	(0.008)
β_3: $T_{i,t}\, T_{j,t}\, g_{j,t}$					-0.011	-0.015**	0.001	0.008
					(0.010)	(0.006)	(0.006)	(0.005)
δ_0: $\pi^{e,h}_{j,t}$	-0.001	0.004	-0.012	0.008	-0.007	0.010	-0.018	0.016
	(0.019)	(0.017)	(0.029)	(0.027)	(0.027)	(0.019)	(0.030)	(0.028)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.026*	0.032	0.060**	0.036*	0.034*	0.024	0.069**	0.032*
	(0.014)	(0.020)	(0.029)	(0.021)	(0.017)	(0.017)	(0.027)	(0.018)
δ_2: $T_{i,t}\, \pi^{e,h}_{j,t}$					0.031	0.026	0.038	-0.000
					(0.032)	(0.026)	(0.030)	(0.031)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					-0.034*	-0.027	-0.044*	-0.003
					(0.019)	(0.020)	(0.022)	(0.016)
Panel B: Downstream signals ($g_{j,t}$ = customer purchase growth $g^{r}_{j,t}$)								
β_0: $g_{j,t}$	-0.001	-0.001	0.000	-0.000	0.000	0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
β_1: $T_{j,t}\, g_{j,t}$	0.002	0.001	-0.000	0.002**	0.001	-0.000	-0.002	0.001
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
β_2: $T_{i,t}\, g_{j,t}$					-0.000	-0.000	-0.001	-0.002
					(0.001)	(0.002)	(0.001)	(0.001)
β_3: $T_{i,t}\, T_{j,t}\, g_{j,t}$					0.001	0.003	0.004	0.001
					(0.001)	(0.003)	(0.003)	(0.002)
δ_0: $\pi^{e,h}_{j,t}$	-0.030	-0.028	-0.046	-0.010	-0.020	-0.036	-0.033	-0.008
	(0.024)	(0.023)	(0.031)	(0.027)	(0.014)	(0.021)	(0.021)	(0.024)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.079***	0.085***	0.113***	0.064**	0.056***	0.098***	0.089**	0.050
	(0.021)	(0.024)	(0.030)	(0.025)	(0.017)	(0.032)	(0.032)	(0.033)
δ_2: $T_{i,t}\, \pi^{e,h}_{j,t}$					0.054**	0.077**	0.071**	0.042
					(0.019)	(0.033)	(0.031)	(0.034)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					-0.066**	-0.098**	-0.066	-0.053
					(0.025)	(0.046)	(0.042)	(0.038)
ρ: $\pi^{e,h}_{i,t-1}$	0.248***	0.262***	0.261***	0.261***	0.281***	0.290***	0.289***	0.289***
	(0.061)	(0.053)	(0.053)	(0.053)	(0.046)	(0.045)	(0.045)	(0.045)
Observations	166,370	166,370	166,370	166,370	497,406	497,406	497,406	497,406
R^2	0.062	0.069	0.070	0.069	0.078	0.083	0.083	0.083

Notes: Both panels report coefficients from the joint pooled regression in equation (6), estimated on the set of dyads (i,j,t) with $j \in \mathcal{N}_{i,t}$. The dependent variable is firm i's inflation expectation $\pi^{e,h}_{i,t}$ pooled across $h \in \{6,12\}$. Panel A shows upstream signals where $g_{j,t}$ is supplier j's year-on-year own price growth; Panel B shows downstream signals where $g_{j,t}$ is customer j's year-on-year purchase growth. In both panels $\pi^{e,h}_{j,t}$ is partner j's pooled inflation expectation. Columns (1)–(4) report the specification on untreated firms only ($T_{i,t}=0$; $\beta_2, \beta_3, \delta_2, \delta_3$ suppressed); columns (5)–(8) report the full-sample specification. Weighting schemes: *unif* (uniform; baseline), *brd* (breadth-weighted), *val* (transaction-value), *shr* (partner-sales-share tilted). All specifications include firm, horizon, sector×quarter, size×quarter, area×quarter, and direction×treatment-pair×quarter fixed effects (the last in full-sample columns only). Standard errors use Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

belief-based learning and price-based inference varies systematically across relationships, and this heterogeneity is driven primarily by supplier-side rather than buyer-side characteristics. We turn to that evidence next.

4 Supplier heterogeneity in information transmission

This section presents the central empirical fact of the paper. Upstream information transmission is sharply heterogeneous across relationships. Belief-based learning is strongest in economically important and persistent supplier relationships, while price-based learning is strongest when suppliers have broader networks. The same gradients weaken sharply, and in some cases reverse, when the relevant margins are measured from the buyer's side rather than the supplier's side. By contrast, downstream learning from customers operates through the belief channel on average but exhibits little systematic heterogeneity.

These patterns are difficult to reconcile with buyer-side theories in which firms choose how much attention to allocate to partner signals. If buyer attention were the operative margin, the gradients should be at least as strong when relationship importance is measured from the buyer's perspective. Instead, the data point to *supplier-originated disclosure*: suppliers communicate more in durable and important relationships, while buyers rely more on prices when disclosure is harder to sustain. This supplier-side asymmetry is the core empirical motivation for the model developed in Section 5.

We study three margins of heterogeneity. The first two capture relationship importance from the partner's side: a value-based share and a persistence-based share. The third captures partner breadth, measured by the supplier's direction-specific number of links. We then compare these partner-side gradients with the corresponding own-side definitions, which measure the importance of partner j in firm i's transactions or firm i's own network position. Throughout, we use the pooled-direction design introduced in Section 3.2, pool horizons $h \in \{6, 12\}$, restrict attention to untreated firms ($T_{i,t} = 0$), and use uniform within-firm link weights.

4.1 Partner-side gradients

We augment equation (6) by interacting each signal and its $T_{j,t}$ interaction with a centered log heterogeneity variable $\widetilde{z}_{ij,t}$:

$$\begin{aligned}\pi^{e,h}_{i,t} = {} & \alpha_i + \alpha_h + \alpha_{s(i)\times t} + \alpha_{d(i)\times t} + \alpha_{a(i)\times t} + \alpha_{T_{j,t}\times t} + \rho\, \pi^{e,h}_{i,t-1} + \gamma_1 \widetilde{z}_{ij,t} + \gamma_2 T_{j,t}\, \widetilde{z}_{ij,t} \\ & + \Big[\beta_0 + \beta_1 T_{j,t} + \beta_2 \widetilde{z}_{ij,t} + \beta_3 (T_{j,t} \times \widetilde{z}_{ij,t})\Big]\, g_{j,t} \\ & + \Big[\delta_0 + \delta_1 T_{j,t} + \delta_2 \widetilde{z}_{ij,t} + \delta_3 (T_{j,t} \times \widetilde{z}_{ij,t})\Big]\, \pi^{e,h}_{j,t} + \varepsilon_{ij,t}, \end{aligned} \tag{7}$$

where the lower-order terms absorb direct effects of relationship characteristics on expectations, so that β_3 and δ_3 capture differential *signal loading* rather than level shifts. The variable $\widetilde{z}_{ij,t}$ is the centered log of partner-side value share (column 1), persistence share (column 2), or partner breadth (column 3); column (4) enters value and persistence jointly. All specifications use uniform within-firm link weights and Driscoll–Kraay standard errors (bandwidth 2).

Table 4 reports the coefficient estimates, and Table 5 translates them into treated-partner premia at P10 and P90 of each margin. The sharpest heterogeneity arises upstream, reported in Panel A, and that is where we focus.

The two partner-side size margins tell a consistent story. Belief-based learning rises significantly with both value share ($\delta_3 = 0.011$) and relationship persistence ($\delta_3 = 0.035$). The persistence gradient is substantially larger. When value and persistence are entered jointly in column (4), the value gradient attenuates to insignificance while the persistence gradient remains strong (0.022). Persistence, not contemporaneous trade volume, is therefore the dominant relationship margin. Evaluated from P10 to P90, the persistence gradient spans $\Delta = 0.133$, implying that belief-based learning is effectively absent in transient relationships and concentrated in durable ones.

The price channel moves in the opposite direction. Along the persistence margin, the upstream treated-partner premium in supplier prices declines from 0.023 at P10 to -0.018 at P90, for a total gradient of -0.040. Thus, in durable supplier relationships, direct communication strengthens while price-based inference fades. This is exactly the substitution pattern one would expect if persistent relationships make disclosure easier to sustain.

Partner breadth reveals an even sharper crossover. The upstream belief premium declines significantly with breadth ($\delta_3 = -0.022$), while the upstream price premium increases ($\beta_3 = 0.011$). At P10 of breadth, the belief premium is 0.083 and the price premium is -0.025; by P90, the belief premium falls to an insignificant -0.003, while the price premium rises to 0.020. As supplier networks become broader, prices replace beliefs as the main vehicle of information transmission. This pattern is consistent with broad-network suppliers being harder to communicate with intensively on any one link, making buyers rely more on price-based inference. Appendix P shows that nonparametric bin regressions display the same crossover.

Downstream heterogeneity is much weaker. Belief premia on customer links are positive and stable (0.071–0.077) across all partner-side margins, while the gradients are small and statistically insignificant. The transaction-based downstream channel remains economically small throughout. Thus, downstream learning operates through the belief channel on average, but without the structured heterogeneity visible upstream.

Table 5 translates these estimates into treated-partner premia at P10 and P90 of each partner-side margin. The crossover between beliefs and prices is especially clear for persistence and breadth. Along persistence, the belief premium rises steeply from P10 to P90 while the price premium falls. Along breadth, the belief premium declines while the price premium rises. Appendix M, Figure M.1, shows the same patterns graphically.

Partner-side versus own-side heterogeneity

The preceding results measure relationship importance from the partner's perspective. Appendix M reports the analogous estimates under own-side definitions, which instead measure the importance of partner j in firm i's transactions or firm i's own network position. The contrast is sharp: the upstream belief gradient weakens substantially, and in several cases reverses sign, under own-side definitions.

This reversal is most striking for persistence. In the joint value-persistence specification, the upstream belief gradient is 0.090 when persistence is measured from the supplier's side, but -0.030 and insignificant when measured from the buyer's side. The same pattern appears for value-based size and breadth: partner-side definitions generate strong or significant upstream belief gradients, whereas own-side definitions produce flat, weak, or reversed gradients.[15]

This asymmetry is difficult to reconcile with buyer-side theories of attention or information

[15]Full own-side regression tables and percentile marginals are reported in Appendix M.

Table 4: Heterogeneity in learning from partners

	(1) Value $\widetilde{\ell}_{ij}^{(j),\text{val}}$	(2) Persistence $\widetilde{\ell}_{ij}^{(j),\text{per}}$	(3) Breadth $\widetilde{b}_{j,t}^{(j)}$	(4) Val. + Pers. Joint
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g_{j,t}^{p}$)				
β_1: $T_{j,t}\, g_{j,t}$	0.006	0.001	−0.001	0.005
	(0.006)	(0.007)	(0.007)	(0.007)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(j),\text{val}}$	−0.003*			−0.000
	(0.001)			(0.002)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		−0.011**		−0.007
		(0.005)		(0.006)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{b}_{j,t}^{(j)}$			0.011**	
			(0.005)	
δ_1: $T_{j,t}\, \pi_{j,t}^{e,h}$	0.033**	0.042***	0.037**	0.037**
	(0.014)	(0.014)	(0.015)	(0.015)
δ_3: $T_{j,t}\, \pi_{j,t}^{e,h} \times \widetilde{\ell}_{ij}^{(j),\text{val}}$	0.011***			0.004
	(0.003)			(0.004)
δ_3: $T_{j,t}\, \pi_{j,t}^{e,h} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		0.035***		0.022***
		(0.006)		(0.006)
δ_3: $T_{j,t}\, \pi_{j,t}^{e,h} \times \widetilde{b}_{j,t}^{(j)}$			−0.022**	
			(0.010)	
Panel B: Downstream signals ($g_{j,t}$ = customer purchase growth $g_{j,t}^{r}$)				
β_1: $T_{j,t}\, g_{j,t}$	0.002	0.002	0.002	0.002
	(0.002)	(0.002)	(0.002)	(0.002)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(j),\text{val}}$	−0.000			0.000
	(0.001)			(0.001)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		−0.000		−0.001
		(0.001)		(0.001)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{b}_{j,t}^{(j)}$			−0.000	
			(0.000)	
δ_1: $T_{j,t}\, \pi_{j,t}^{e,h}$	0.071***	0.075***	0.077***	0.073***
	(0.018)	(0.019)	(0.019)	(0.018)
δ_3: $T_{j,t}\, \pi_{j,t}^{e,h} \times \widetilde{\ell}_{ij}^{(j),\text{val}}$	0.003			0.004
	(0.003)			(0.004)
δ_3: $T_{j,t}\, \pi_{j,t}^{e,h} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		0.003		−0.003
		(0.007)		(0.008)
δ_3: $T_{j,t}\, \pi_{j,t}^{e,h} \times \widetilde{b}_{j,t}^{(j)}$			−0.003	
			(0.003)	
ρ: $\pi_{i,t-1}^{e,h}$	0.248***	0.248***	0.248***	0.248***
	(0.061)	(0.061)	(0.061)	(0.061)
Observations	166,370	166,370	166,146	166,370
R^2	0.066	0.067	0.066	0.067

Notes: Each column estimates equation (7) on untreated firms ($T_{i,t} = 0$) with uniform link weights and pooled horizons $h \in \{6, 12\}$. The heterogeneity variable is: column (1) the centered log of firm i's share in partner j's total transaction value cumulated over the sample; column (2) the centered log of the dyad's persistence share in partner j's total relationship persistence; column (3) the centered log of partner j's direction-specific number of trading partners; column (4) value and persistence entered jointly. Since both panels are extracted from a single pooled-direction regression, ρ and fit statistics are common across directions. For brevity, only $T_{j,t}$ interaction terms and their heterogeneity gradients are shown.

Table 5: Treated-partner learning premia at P10 and P90 (partner-side)

	P10	P90	Δ(P90–P10)
Panel A: Value-based size $\widetilde{\ell}_{ij}^{(j),\mathrm{val}}$			
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)	−0.021	0.081***	0.101***
	(0.020)	(0.019)	(0.028)
$T_{j,t} \times g_{j,t}^{p}$	0.018	−0.004	−0.022*
	(0.011)	(0.006)	(0.012)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.057**	0.084***	0.027
	(0.021)	(0.024)	(0.029)
$T_{j,t} \times g_{j,t}^{r}$	0.002	0.001	−0.001
	(0.004)	(0.002)	(0.005)
Panel B: Persistence $\widetilde{\ell}_{ij}^{(j),\mathrm{per}}$			
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)	−0.028	0.105***	0.133***
	(0.020)	(0.017)	(0.023)
$T_{j,t} \times g_{j,t}^{p}$	0.023*	−0.018	−0.040**
	(0.011)	(0.012)	(0.019)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.069**	0.079***	0.010
	(0.027)	(0.020)	(0.027)
$T_{j,t} \times g_{j,t}^{r}$	0.002	0.001	−0.001
	(0.003)	(0.002)	(0.004)
Panel C: Partner breadth $\widetilde{b}_{j,t}^{(j)}$			
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)	0.083***	−0.003	−0.086**
	(0.024)	(0.024)	(0.038)
$T_{j,t} \times g_{j,t}^{p}$	−0.025**	0.020	0.045**
	(0.012)	(0.013)	(0.021)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.085***	0.071***	−0.014
	(0.021)	(0.019)	(0.011)
$T_{j,t} \times g_{j,t}^{r}$	0.002	0.002	−0.000
	(0.002)	(0.002)	(0.002)

Notes: Entries evaluate the treated-partner premium at the 10th and 90th percentiles of the centered log heterogeneity variable, from columns (1), (2), and (3) of Table 4, respectively. Standard errors by the delta method. *, **, *** denote significance at the 10%, 5%, and 1% levels.

acquisition. If buyers were choosing how much to learn from suppliers, the gradient should be at least as strong when relationship importance is measured from the buyer's perspective. Instead, the data indicate that upstream belief transmission depends primarily on supplier-side relationship characteristics. Suppliers communicate more with buyers who account for a larger share of their sales and with whom they expect the relationship to persist. This pattern is more naturally consistent with supplier-originated disclosure.

The heterogeneity analysis therefore delivers two connected facts. First, upstream learning shifts systematically between beliefs and prices across supplier relationships: belief-based learning is strongest in durable and important links, while price-based inference is strongest for broad-network suppliers. Second, these gradients are fundamentally partner-side rather than own-side. Together, these patterns are difficult to reconcile with broad buyer-driven explanations and provide the empirical foundation for the disclosure model developed in the next section. Downstream learning from customers remains present on average but shows little structured heterogeneity across any of these margins (Appendix Q).

State-dependent transmission

We close this section with an additional empirical regularity: belief-based transmission is stronger when aggregate nominal uncertainty is high. We document this pattern here as part of the reduced-form characterization of information transmission in the data. This evidence is supplementary to the cross-sectional heterogeneity documented above and helps characterize the environments in which partner beliefs matter more for expectation formation.

Empirical design. We augment the baseline expectation-formation specification (6) by interacting every partner signal with demeaned squared CPI inflation $\tilde{\pi}_t^2 \equiv \pi_t^2 - \overline{\pi^2}$, a proxy for aggregate nominal uncertainty, and standardize this proxy by its cross-quarter standard deviation so that coefficients measure the effect of a one-standard-deviation increase.[16] In the model, the variance of the common co amponent $\sigma_{q,j}^2$ drives the cooperative surplus from disclosure; empirically, π_t^2 is a natural proxy because the cross-sectional dispersion of cost shocks rises monotonically with the level of aggregate inflation during our sample.[17] The dependent variable is firm i's inflation expectation $\pi_{i,t}^{e,h}$ (percentage points); the full estimating equation is in Appendix R.1. The key coefficients are the belief-channel interactions with uncertainty:

$$T_j \times \pi_{j,t}^e \times \tilde{\pi}_t^2 \qquad \text{and} \qquad T_i \times T_j \times \pi_{j,t}^e \times \tilde{\pi}_t^2,$$

where the second term is included in the full sample to capture how the marginal effect of uncertainty differs on links where both parties are treated.

Table 6 reports the estimates. Column (1) restricts to untreated buyers. The downstream belief interaction is positive and highly significant: a one-standard-deviation increase in inflation uncertainty raises the treated-partner belief premium by 0.047 p.p., roughly a doubling relative to the baseline downstream premium of 0.085. The upstream belief interaction is also positive and significant (0.027). Column (2) adds the full sample with own-treatment interactions.

[16] Demeaning ensures that the baseline treated-partner premia are evaluated at average uncertainty, while the interaction coefficients capture how transmission slopes with deviations from that average.

[17] Squared CPI inflation has the advantage of being directly observable at the quarterly frequency and exogenous to any individual firm's behavior. The proxy is also consistent with evidence that price dispersion rises with the level of inflation; see, e.g., Coibion et al. (2020).

The treated-supplier coefficients are positive, and the own-treatment interaction is negative and highly significant in both directions (-0.050 to -0.059): when both parties are treated, disclosure is already more likely on a larger set of links, so the marginal effect of higher uncertainty is smaller.

Table 6: State-dependent learning

	Dep. var.: $\pi^{e,h}_{i,t}$ (p.p.)	
	(1) Untreated	(2) All
Belief channel $\times\tilde{\pi}^2_t$/sd		
δ^{π}_1: $T_j \times \pi^e_{j,t} \times \tilde{\pi}^2_t$ (up)	0.027^{**}	0.051^{***}
	(0.012)	(0.016)
δ^{π}_1: $T_j \times \pi^e_{j,t} \times \tilde{\pi}^2_t$ (down)	0.047^{***}	0.032^{**}
	(0.014)	(0.013)
δ^{π}_3: $T_i \times T_j \times \pi^e_{j,t} \times \tilde{\pi}^2_t$ (up)		-0.059^{***}
		(0.015)
δ^{π}_3: $T_i \times T_j \times \pi^e_{j,t} \times \tilde{\pi}^2_t$ (down)		-0.050^{***}
		(0.013)
Observations	166,370	497,406
R^2	0.063	0.079

Notes: Dependent variable is firm i's inflation expectation $\pi^{e,h}_{i,t}$ (percentage points). The specification augments equation (6) by interacting all partner signals with $\tilde{\pi}^2_t/\text{sd}(\tilde{\pi}^2_t)$, demeaned squared CPI inflation standardized by its cross-quarter standard deviation (full equation in Appendix R.1). Demeaning implies that baseline treated-partner premia are evaluated at average uncertainty, while the interaction coefficients capture how premia change with deviations from average uncertainty. Coefficients measure the change in the belief premium per one-standard-deviation increase in inflation uncertainty. Only belief-channel interactions are shown; price-channel interactions and lower-order terms are included but not reported. Column (1) restricts to untreated buyers ($T_i = 0$); column (2) includes all buyers and adds own-treatment interactions. All specifications include firm, horizon, sector×quarter, size×quarter, area×quarter, and T_j×quarter fixed effects; column (2) replaces T_j×quarter with treatment-pair×quarter fixed effects. Standard errors use Driscoll–Kraay correction (bandwidth 2). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5 Model

The combination of persistent experimental variation and the universe of bilateral transactions allows us to separately identify the belief and price channels of information transmission and to measure the trade-volume distortion from price-based inference directly. The equilibrium of the model developed here connects this distortion to a welfare cost in supply chains, quantified in Section 6. The organizing principle is a directional asymmetry: the value of upstream information transmission is *first order* in the informational content of prices, whereas the value of downstream transmission is at most *second order* (Proposition 1). This asymmetry determines which direction of information flow can justify costly relational enforcement, and thereby which direction exhibits cross-link heterogeneity tied to relationship characteristics.

We proceed in five steps. Section 5.1 sets up the production network: each buyer is matched to a single supplier per upstream industry (Figure 3), so that bilateral prices confound the supplier's common cost component with a match-specific idiosyncratic term. Section 5.2 solves the bilateral pricing problem: under dispersed information, the buyer's use of the price as a signal about the supplier's costs reduces the perceived demand elasticity and raises the markup above the frictionless level, creating a first-order/second-order asymmetry in the value of information

transmission (Proposition 1). Section 5.3 derives the link-level incentive constraint that governs supplier disclosure (Proposition 2): disclosure is sustainable only on links where the relationship is sufficiently durable and valuable, generating cross-link heterogeneity in upstream belief transmission. Appendix F embeds the mechanism in the full production network, introducing passive attention signals that provide a baseline information channel on all links regardless of relationship characteristics—accounting for the lack of structured heterogeneity in downstream transmission—and yielding a system of linear prices and beliefs with an endogenously determined disclosure set $\mathcal{R}_t$. Section 6 takes the model's predictions to the data.

The general equilibrium closure—household preferences, monetary policy, and market clearing—follows La'O and Tahbaz-Salehi (2022) and is in Appendix H. The key simplification is that monetary policy pins down the nominal wage, so individual firms treat it as exogenous; the bilateral results (Propositions 1 and 2) and their empirical predictions are therefore independent of the closure. Full matrix derivations are in Appendix F.

5.1 The firm-to-firm production network

Production and matching. We embed the bilateral information transmission problem in a production network pricing framework following La'O and Tahbaz-Salehi (2022). There are n industries connected by an input-output matrix $A = [a_{ij}]$, where a_{ij} is the share of industry j in the production technology of industry i. Industry i has labor share $\alpha_i = 1 - \sum_j a_{ij}$. Each industry contains a continuum of firms indexed by $k \in [0, 1]$.

Firm k in industry i produces output using a Cobb–Douglas technology:

$$y_{ik,t} = z_{ik,t}\, l_{ik,t}^{\alpha_i} \prod_{j=1}^{n} m_{ijk,t}^{a_{ij}}, \tag{8}$$

where $l_{ik,t}$ is labor hired at the nominal wage w_t, $m_{ijk,t}$ is the intermediate input purchased from the firm's supplier in industry j, and $z_{ik,t}$ is firm-specific productivity.

Within each industry, firms are matched to specific trading partners upstream and downstream. For each upstream industry j with $a_{ij} > 0$, firm k in industry i is matched to a *single* supplier $\ell_j(k)$ in industry j, who charges a bilateral input price $p_{\ell_j(k),t}$ specific to the match. The identity of $\ell_j(k)$ differs across buyers k, so distinct firms in industry i face distinct bilateral prices. We write $\mathcal{N}_{ik}^{\text{up}}$ for the set of upstream suppliers of firm k and $\mathcal{N}_{ik}^{\text{dn}}$ for its set of downstream buyers.[18] Each match survives to the next period with probability $\rho_{ij} \in (0, 1)$; upon separation, the buyer draws a new supplier at random from the same industry. Match turnover thus changes the identity of k's upstream partners (when k's own matches break) and of k's downstream partners (when k's customers draw new suppliers). All firms discount future profits at rate $\beta \in (0, 1)$. Figure 3 illustrates an example of the matching structure.

The bilateral matching structure is the key departure from La'O and Tahbaz-Salehi (2022), where firms purchase from industries and all learning occurs through industry-level prices. Here, because firm k's suppliers are specific to k, the bilateral prices and any disclosed beliefs that k

[18]The one-supplier-per-upstream-industry restriction is the simplest matching structure that generates the bilateral information friction. In principle, one could allow each buyer to source from multiple suppliers producing differentiated products within the same industry. The information friction and the incentive to disclose would operate similarly on each link so long as the number of matches is finite, but suppliers would then interact strategically in pricing and disclosure decisions toward the same buyer, complicating the analysis considerably; see Section 7 for further discussion.

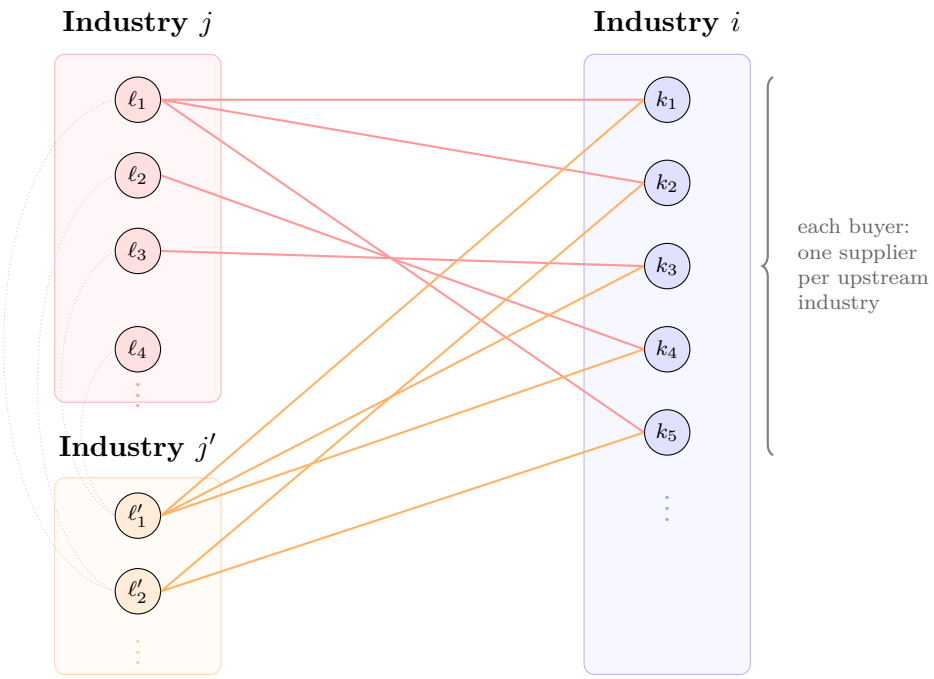


Figure 3: Bilateral matching. Each buyer in industry i draws one supplier per upstream industry (j and j'); draws are independent across buyers. Supplier degree varies: ℓ_1 serves three buyers in i, ℓ_3 one, ℓ_4 none. Dotted curves show that j and j' also trade with each other. Each link carries a match-specific bilateral price.

observes constitute *private* information channels unavailable to other firms in the same industry. The random matching of buyers to suppliers plays the same role as the random assignment of consumers to islands in Lorenzoni (2009): randomness generates match-specific noise, and the inability to aggregate across trading partners sustains the information friction (see Section 5.2).

Matching also gives each supplier a degree of monopoly power over its matched buyers. To limit this power, we follow a standard device in the search-and-matching literature: each industry j has a competitive *wholesale producer* that sells at the industry-average marginal cost $Q_{j,t}$, which aggregates over all firms in j so that firm-specific cost variation washes out and $Q_{j,t}$ depends only on the industry-common component of marginal cost (defined in the next paragraph). A buyer can switch to the wholesale producer at an iceberg cost—it must purchase $\tau \geq 1$ units to obtain one usable unit, where τ is drawn from a Pareto distribution with shape parameter $\varepsilon_i - 1$, and $\varepsilon_i > 1$ is the bilateral demand elasticity faced by suppliers selling to industry i. The matched supplier therefore faces a downward-sloping demand curve: it can charge above $Q_{j,t}$, but at the risk of losing the buyer to the wholesale alternative. The Demand paragraph below derives the resulting constant-elasticity demand schedule.

We decompose firm-level productivity $z_{ik,t}$ into a component common to all firms in industry i and a firm-specific residual: $\log z_{ik,t} = e_{i,t} - z_{k,t}$, where $e_{i,t} \sim \mathcal{N}(0, \sigma^2_{e,i})$ is an industry-common productivity shock—capturing sector-specific technology or cost conditions such as factor-market disruptions—and $z_{k,t} \sim \mathcal{N}(0, \sigma^2_{z,i})$ is a firm-specific productivity shock, mutually independent across firms and independent of all aggregate variables. We place no restriction on the correlation structure of $\{e_{i,t}\}_{i=1}^n$ across industries: common shocks are in general correlated because they propagate through the production network.

The log nominal marginal cost of firm k in industry i is, to a first-order approximation,[19]

$$\log mc_{ik,t} \;=\; \alpha_i \log w_t \;+\; \underbrace{e_{i,t}}_{\text{industry-common}} \;-\; \underbrace{z_{k,t}}_{\text{firm-specific}} \;+\; \sum_{j=1}^{n} a_{ij} \log p_{\ell_j(k),t}. \tag{9}$$

Define the *common component* $q_{i,t}$ for industry i recursively as

$$q_{i,t} \;\equiv\; \alpha_i(\log w_t - \bar{w}) \;+\; e_{i,t} \;+\; \sum_{j=1}^{n} a_{ij}\, q_{j,t}. \tag{10}$$

The object $q_{i,t}$ collects all within-industry correlated variation in marginal cost: industry i's own common shock, the aggregate wage, and the common shocks of upstream industries propagated through the input-output structure. It is common to all firms in industry i and does not depend on k. Because $q_{i,t}$ is a linear combination of jointly Gaussian shocks, the vector $\mathbf{q}_t \equiv (q_{1,t}, \ldots, q_{n,t})'$ is jointly Gaussian with mean zero and covariance matrix $\Sigma_q \equiv \mathrm{Var}(\mathbf{q}_t)$. The diagonal elements $\sigma^2_{q,i} \equiv \mathrm{Var}(q_{i,t})$ are industry-specific; the off-diagonal elements $\mathrm{Cov}(q_{i,t}, q_{j,t})$ are in general nonzero because common shocks propagate across industries through the input-output structure.[20]

The log marginal cost of firm k in industry i can therefore be decomposed as

$$\log mc_{ik,t} \;=\; \overline{mc}_i \;+\; q_{i,t} \;+\; \zeta_{ik,t}, \tag{11}$$

where $\overline{mc}_i$ is the time-invariant steady-state log marginal cost, common to all firms in industry i and known from the history of trading relationships, and $\zeta_{ik,t}$ is the firm-specific idiosyncratic residual:[21]

$$\zeta_{ik,t} \;\equiv\; -z_{k,t} \;+\; \sum_{j=1}^{n} a_{ij}\Big(\log p_{\ell_j(k),t} - \bar{p}_{ij} - q_{j,t}\Big), \tag{12}$$

where $\bar{p}_{ij} \equiv \log \mu_i^* + \overline{mc}_j$ is the steady-state bilateral price on any link from industry j to industry i, and $\mu_i^* \equiv \varepsilon_i/(\varepsilon_i - 1)$ is the frictionless markup (derived in Section 5.2). Because $\bar{p}_{ij}$ is defined at the frictionless benchmark μ_i^*, the residual $\zeta_{ik,t}$ absorbs both the stochastic firm-specific variation and, on non-disclosure links, the deterministic markup deviation $(\log \mu^D_{ij,t} - \log \mu_i^*)$ derived in Section 5.2.[22] Conditional on the disclosure configuration, the stochastic component of $\zeta_{ik,t}$ is Gaussian with mean zero and variance $\sigma^2_{\zeta,i}$ that depends on the network position of industry i.

[19] All expressions in this section are first order in the perturbation parameter δ that indexes the distance from the symmetric steady state; we suppress the remainder $o(\delta)$ throughout. The marginal cost is conditional on the buyer sourcing from the matched supplier; to first order, the switching probability is constant across firms within each industry pair and absorbed into the steady-state level $\overline{mc}_i$. The decomposition follows La'O and Tahbaz-Salehi (2022) (Lemma 2).

[20] Iterating (10) yields the closed form $q_{i,t} = \sum_j h_{ij}\alpha_j(\log w_t - \bar{w}) + \sum_j h_{ij}\, e_{j,t}$, where h_{ij} is the (i,j) element of the Leontief inverse $H = (I - A)^{-1}$. The covariance matrix is $\Sigma_q = H\,\Sigma_e\,H'$, where Σ_e is the covariance matrix of $\{e_{i,t}\}$.

[21] The steady-state marginal cost $\overline{mc}_i = \alpha_i \bar{w} + \sum_j a_{ij}(\overline{mc}_j + \log \mu_i^*)$ depends only on industry primitives. In steady state, the markup that any supplier in j charges to any buyer in i is μ_i^* (determined by the buyer's demand elasticity), and all firms in j share the same marginal cost $\overline{mc}_j$.

[22] The markup deviation depends on the Kalman weight $\omega_{ij,t}$, which is a function of signal precisions and the disclosure configuration—not of shock realizations. It is therefore common to all firms in the (i,j) industry pair and known to both parties. Both sides subtract it when updating beliefs, so it does not affect signal extraction.

Information structure. Each firm k in industry i has an information set $\mathcal{I}_{k,t}$ that evolves through the stages of each period. At the start of each period, the firm has private signals on two sets of objects:

(a) **Idiosyncratic shocks.** Firm k observes its own productivity shock $z_{k,t}$. For each upstream partner $\ell \in \mathcal{N}_{ik}^{\text{up}}$ in industry j, firm k receives a noisy private signal

$$s_{z,k\ell,t} = z_{\ell,t} + \vartheta_{k\ell,t}, \qquad \vartheta_{k\ell,t} \sim \mathcal{N}(0, \sigma^2_{\vartheta,j}),$$

independent across firms and links, as in the signal structure of La'O and Tahbaz-Salehi (2022) adapted to the firm-specific decomposition.[23] Because firm k does not observe upstream bilateral prices before setting its own prices (see the timing protocol below), these signals—together with $z_{k,t}$—are k's only source of information about the idiosyncratic shocks of its suppliers. They allow k to form the prior $\hat{\zeta}_{ik,t}$ on the idiosyncratic residual (12) before entering bilateral negotiations. Since $\hat{\zeta}_{ik,t}$ depends on k's match-specific upstream draws, it is private to k and independent across firms within industry i.

(b) **Common component.** Firm k receives a noisy private signal about its own industry's common component:

$$s_{q,k,t} = q_{i,t} + \nu_{k,t}, \qquad \nu_{k,t} \sim \mathcal{N}(0, \sigma^2_{\nu,i}),$$

independent across firms.

The two sets of signals serve distinct roles. The idiosyncratic signals in (a) resolve uncertainty about $\zeta_{ik,t}$, which is firm-specific and averages out at the industry level. The private signal in (b) is the firm's only *ex ante* source of information about the common component $q_{i,t}$, which drives aggregate pricing distortions. During bilateral negotiations, the firm may acquire additional information about $q_{j,t}$ through disclosed supplier beliefs or the informational content of negotiated prices (Section 5.3); the idiosyncratic prior $\hat{\zeta}_{ik,t}$ is not updated thereafter.[24]

Lemma 1 (Zero overlap)**.** *Under independent matching, any buyer–seller pair (k, ℓ) with k in industry i and ℓ in industry j satisfies:* (i) *$\zeta_{ik,t}$ and $\zeta_{j\ell,t}$ are independent conditional on $\mathbf{q}_t$;* (ii) *the information sets of k and ℓ overlap only in their signals about the common components $q_{i,t}$ and $q_{j,t}$.*

Both parts follow from the independence of $\{z_{k,t}\}$ across firms and the fact that, with probability one, no two firms share a common trading partner under independent matching; the formal proof is in Appendix C. Zero overlap makes information transmission a self-contained bilateral problem, ensures that the price residual on each link (equation (18) below) is a signal about the seller's posterior with independent additive noise, and pins down the outside option upon separation.

[23] In La'O and Tahbaz-Salehi (2022), each firm receives signals about industry-level productivity shocks. Here, productivity decomposes into an industry-common component $e_{i,t}$ (absorbed into $q_{i,t}$) and a firm-specific residual $z_{k,t}$, so the relevant signals are about firm-specific shocks. The formal structure is otherwise identical: Gaussian signals, independent noise, and a common prior.

[24] Appendix G introduces attention signals—noisy observations of trading partners' beliefs generated by passive monitoring of observable activity on every link—when the industry-level belief system requires specifying all information channels. These signals add a baseline channel of information transmission that is insensitive to relationship characteristics, matching the downstream empirical patterns. The core bilateral mechanism does not depend on them. Appendix G endogenizes the signal precision through a costly attention problem.

Timing. Within each period t, events unfold in four stages:

(i) **Shocks and signals.** Nature draws $\{q_{i,t}\}$, $\{z_{k,t}\}$, and signal noise. Each firm k observes $z_{k,t}$, the upstream signals $\{s_{z,k\ell,t}\}$, and the private signal $s_{q,k,t}$, and forms the idiosyncratic prior $\hat{\zeta}_{ik,t}$. Trading partners do *not* observe each other's idiosyncratic components.

(ii) **Disclosure.** Let $\mathcal{R}_t \subseteq \{(k,\ell) : a_{ij} > 0\}$ denote the set of links on which voluntary disclosure occurs—the *disclosure set*, determined endogenously below. On each link $(k,\ell) \in \mathcal{R}_t$, the supplier sends a message $m_{k\ell,t}$ about its belief to the buyer. On links $(k,\ell) \notin \mathcal{R}_t$, no voluntary communication occurs.

(iii) **Pricing and transactions.** All firms simultaneously set bilateral prices using their beginning-of-period information: the private signal $s_{q,k,t}$ and the idiosyncratic prior $\hat{\zeta}_{ik,t}$ from stage (i). Firm k does not observe upstream input prices, which are set simultaneously as in La'O and Tahbaz-Salehi (2022). Buyers then observe posted prices and any disclosed messages from stage (ii), form beliefs about supplier costs, choose input quantities, and transactions occur.

(iv) **End-of-period updating.** After transactions, each firm aggregates all signals observed during the period—its own signal, supplier prices, and any disclosed messages—into a posterior belief. Under the i.i.d. assumption on $q_{i,t}$, this posterior is not payoff-relevant: next period's shock is independent, so accumulated information does not carry forward. With persistent shocks, the posterior would become the prior for period $t+1$, and bilateral learning would affect future pricing through improved priors. The i.i.d. case isolates the within-period demand channel of information transmission from the dynamic pricing channel.

Demand. The wholesale producer's log price is $\log Q_{j,t} = \bar{Q}_j + q_{j,t}$, where $\bar{Q}_j$ is a steady-state constant, so residual uncertainty about $q_{j,t}$ translates directly into uncertainty about the outside option. The buyer can switch to the wholesale producer at an iceberg cost τ drawn from a Pareto distribution with shape parameter $\varepsilon_i - 1$; the supplier does not observe τ. Averaging over the staying probability and combining with Cobb–Douglas expenditure shares yields a constant-elasticity demand schedule (Appendix D):

$$x_{kj,t}(p) \;=\; \left(\frac{p}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}\right)^{-\varepsilon_i} A_{kj,t}, \tag{13}$$

where $A_{kj,t} \equiv a_{ij}\, S_{i,t}/\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ (recall $k \in i$, $S_{i,t}$ the buyer's total nominal input expenditure) is a demand shifter that depends on the buyer's industry and information set but not on the supplier's identity.[25] Because $Q_{j,t}$ is not directly observed, the conditional expectation in (13) depends on the buyer's information. On *disclosure links*, the buyer learns the supplier's posterior $\hat{q}_{j\ell,t} \equiv \mathbf{E}_{j\ell}[q_{j,t}]$ directly, so the demand elasticity is constant at ε_i. On *non-disclosure links*, the buyer extracts information from p, reducing the perceived demand elasticity and generating the markup distortion derived next.

[25] The total elasticity ε_i combines a unit intensive-margin elasticity (Cobb–Douglas expenditure shares) with an extensive-margin elasticity of $\varepsilon_i - 1$ (switching probability).

5.2 Bilateral pricing and surplus

We solve the bilateral problem by backward induction through stages (iv) and (iii) of the timing protocol. We first characterize bilateral surplus as a function of the markup, then derive the equilibrium markup under dispersed information. Section 5.3 completes the analysis by determining when disclosure is sustainable (stage ii); the product-market equilibrium and general equilibrium closure are in Appendix F and Appendix H, respectively.

Bilateral surplus. We characterize bilateral surplus at the symmetric steady state (no uncertainty about marginal costs, all firms in each industry charge identical prices). Let $mc_{j\ell,t}$ denote supplier ℓ's nominal marginal cost in levels and define the demand scale $\Gamma_{k\ell,t} \equiv mc_{j\ell,t}^{1-\varepsilon_i} A_{kj,t}$, which collects the components of bilateral revenue that are taken as given by the supplier when choosing its markup $\mu \equiv p/mc_{j\ell,t}$. The supplier's per-period profit on the link is $\Pi^S_{k\ell,t}(\mu) = h(\mu)\,\Gamma_{k\ell,t}$, where $h(\mu) \equiv (\mu-1)\mu^{-\varepsilon_i}$. Maximizing h yields the first-order condition

$$h'(\mu) = \mu^{-\varepsilon_i - 1}\Big[\varepsilon_i - (\varepsilon_i - 1)\mu\Big] = 0 \qquad \Longrightarrow \qquad \mu_i^* = \frac{\varepsilon_i}{\varepsilon_i - 1}, \tag{14}$$

with $h''(\mu_i^*) < 0$, confirming a maximum. The buyer's surplus is $\Pi^B_{k\ell,t}(\mu) = \mu^{1-\varepsilon_i}\Gamma_{k\ell,t}/(\varepsilon_i - 1)$.[26]

Total bilateral surplus $W_{k\ell,t}(\mu) \equiv \Pi^S_{k\ell,t} + \Pi^B_{k\ell,t}$ satisfies

$$W'_{k\ell,t}(\mu) = \varepsilon_i\,\mu^{-\varepsilon_i - 1}(1-\mu)\,\Gamma_{k\ell,t}, \tag{15}$$

so W is maximized at $\mu = 1$ (marginal-cost pricing) and strictly decreasing for $\mu > 1$. In particular, $W'_{k\ell,t}(\mu_i^*) < 0$: any markup above marginal cost destroys total surplus. Because $h'(\mu_i^*) = 0$ by the first-order condition (14), a small markup change has only a second-order effect on the seller's profit but a first-order effect on buyer surplus. This envelope condition is the source of the directional asymmetry derived below.

Pricing under full disclosure. On a link where the buyer's information set $\mathcal{I}_{k,t}$ does not respond to the supplier's price p—either because the supplier has disclosed $\hat{q}_{j\ell,t}$ directly, or because the buyer has sufficiently precise information from other sources—demand (13) has constant elasticity ε_i and the profit-maximizing markup is μ_i^* from (14). The supplier's pricing rule is:

$$\log p^*_{k\ell,t} \;=\; \bar{p}_{ij} \;+\; \hat{q}_{j\ell,t} \;+\; \hat{\zeta}_{j\ell,t}, \tag{16}$$

where $\bar{p}_{ij}$ and $\hat{q}_{j\ell,t}$ are as defined in Section 5.1, and $\hat{\zeta}_{j\ell,t} \equiv \mathbf{E}_{j\ell}[\zeta_{j\ell,t}]$ is the supplier's posterior about the idiosyncratic component of its own marginal cost.[27]

Pricing under dispersed information. On non-disclosure links, the buyer uses the supplier's price to update beliefs about $q_{j,t}$. The conditional expectation $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ in the demand schedule (13) therefore responds to p, and the supplier internalizes this when choosing its markup.

[26] Appendix E derives Π^B from the CES demand integral and verifies the total surplus and its derivative. Away from the steady state, the demand scale $\Gamma_{k\ell,t}$ depends on beliefs and the markup may deviate from μ_i^*; the welfare comparisons below evaluate surplus differences to first order in δ, holding $\Gamma_{k\ell,t}$ at its steady-state value.

[27] The markup depends only on the buyer's industry i, so μ^* is the same for every link into industry i.

To see the mechanism, note that the bilateral log price can be decomposed as

$$\log p_{k\ell,t} \;=\; \bar{p}_{ij} \;+\; (\log \mu_{k\ell,t} - \log \mu_i^*) \;+\; \hat{q}_{j\ell,t} \;+\; \hat{\zeta}_{j\ell,t}, \tag{17}$$

where the supplier's posteriors $\hat{q}_{j\ell,t}$ and $\hat{\zeta}_{j\ell,t}$ are as defined in (16), and $\mu_{k\ell,t}$ is the equilibrium markup. Both $\bar{p}_{ij}$ and the markup deviation $(\log \mu_{k\ell,t} - \log \mu_i^*)$ are deterministic and known to the buyer.[28] The buyer can therefore subtract them, leaving the *price residual*:

$$\log p_{k\ell,t} - \bar{p}_{ij} - (\log \mu_{k\ell,t} - \log \mu_i^*) \;=\; \hat{q}_{j\ell,t} + \hat{\zeta}_{j\ell,t}. \tag{18}$$

The idiosyncratic component $\hat{\zeta}_{j\ell,t}$ is independent of the seller's posterior $\hat{q}_{j\ell,t}$ (by Lemma 1), so the price residual is a signal about $\hat{q}_{j\ell,t}$ with additive independent noise: the bilateral price confounds $q_{j,t}$ with $\hat{\zeta}_{j\ell,t}$, just as local demand confounds aggregate and idiosyncratic shocks in Lorenzoni (2009).[29] By the law of iterated expectations, $\mathbf{E}[q_{j,t} \mid \hat{q}_{j\ell,t}] = \hat{q}_{j\ell,t}$, so learning the seller's posterior is equivalent to learning about $q_{j,t}$.[30] Let $\omega_{ij,t} \in (0,1)$ denote the Kalman weight the buyer places on the price residual:

$$\omega_{ij,t} = \frac{\gamma_{ij,t}}{\gamma_{ij,t} + \sigma^2_{\zeta,j}}, \qquad \gamma_{ij,t} \equiv \mathrm{Var}(\hat{q}_{j\ell,t} \mid \mathcal{I}^{\mathrm{bi},-p}_{k\ell,t}), \tag{19}$$

where $\mathcal{I}^{\mathrm{bi},-p}_{k\ell,t}$ denotes buyer k's information excluding supplier ℓ's bilateral price.[31] The residual uncertainty $\gamma_{ij,t}$ about the seller's posterior depends on both i and j: a seller with more precise private signals has a more variable posterior, raising $\gamma_{ij,t}$ and $\omega_{ij,t}$; a buyer with a more precise upstream attention signal or stronger input–output connections to industry j can better anticipate the seller's posterior, lowering $\gamma_{ij,t}$.[32] In the Gaussian-linear environment, posterior variances depend on signal precisions, not on realizations. Since all buyers in industry i share the same precisions and disclosure configuration, $\gamma_{ij,t}$ is common across firms within the industry pair—hence the subscript (i,j).

Markup distortion. A one-unit increase in $\log p$ shifts the buyer's posterior upward by $\omega_{ij,t}$, which raises $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ and partially offsets the demand reduction from the higher price. From the supplier's perspective, the effective demand elasticity falls from ε_i to $\varepsilon_i(1-\omega_{ij,t})$: a fraction

[28] The equilibrium markup $\mu_{k\ell,t}$ depends on $\omega_{ij,t}$ (derived below), which is determined by signal precisions and the disclosure configuration $\mathcal{R}_t$, not on the realizations of shocks within a period.

[29] The key extension beyond Lorenzoni (2009) is that information transmission is endogenous: on disclosure links, the supplier communicates $\hat{q}_{j\ell,t}$ directly, eliminating the confounding and the associated markup distortion. The model asks when this communication is credible.

[30] In (18) we net out only constants and regime-common-knowledge components; the remaining $\hat{\zeta}_{j\ell,t}$ is orthogonal link-level noise under Lemma 1, so $\omega_{ij,t}$ in (19) summarizes the informativeness of bilateral prices.

[31] We adopt Gaussian signals and Bayesian updating to obtain a tractable expression for the signal-extraction weight $\omega_{ij,t}$ and the first-order surplus comparison in (23). The qualitative mechanism does not require full Bayesian rationality: any belief-formation rule under which buyers place positive weight on supplier prices absent disclosure (so price-based learning affects perceived demand) and on disclosed beliefs when available preserves the directional asymmetry in Proposition 1. Alternative updating rules would change the mapping from observables into $\omega_{ij,t}$ and hence the quantitative magnitudes.

[32] Under a regularity condition on the input-output structure (Assumption 2 in Appendix E), $\omega_{ij,t}$ is *increasing* in supplier breadth. Supplier breadth has competing effects: higher $\sigma^2_{\zeta,j}$ (denominator) pushes ω down, but higher residual uncertainty $\gamma_{ij,t}$ (numerator) pushes it up—a broad supplier's common component $q_{j,t}$ loads on many diverse shocks that the buyer cannot easily learn from its private signal. The assumption ensures the latter dominates.

$\omega_{ij,t}$ of each price increase is absorbed by the belief shift rather than by lost demand. The supplier's problem on non-disclosure links is therefore

$$\max_{\mu} (\mu - 1)\, \mu^{-\varepsilon_i(1-\omega_{ij,t})}\, \widetilde{\Gamma}_{k\ell,t}, \tag{20}$$

where $\widetilde{\Gamma}_{k\ell,t} \equiv mc_{j\ell,t}^{1-\varepsilon_i(1-\omega_{ij,t})}\, A_{kj,t}$ collects terms that do not depend on μ. The first-order condition yields the equilibrium markup

$$\mu_{ij,t}^{D} = \frac{\varepsilon_i(1-\omega_{ij,t})}{\varepsilon_i(1-\omega_{ij,t}) - 1} > \mu_i^*, \tag{21}$$

as in Gaballo and Paciello (2025). The distortion is increasing in $\omega_{ij,t}$: the more the buyer learns from the price, the lower the perceived elasticity and the higher the markup. A first-order expansion gives

$$\mu_{ij,t}^{D} - \mu_i^* = \frac{\varepsilon_i\, \omega_{ij,t}}{(\varepsilon_i - 1)^2} + O\Big(\omega_{ij,t}^2\Big). \tag{22}$$

On disclosure links, $\omega_{ij,t} = 0$ by sufficiency of the disclosed belief, and the markup returns to μ_i^*.

The profit gain from non-disclosure for the seller is

$$\Delta\Pi_{ij,t}^{D} \equiv h_{ij,t}^{D} - h(\mu_i^*) = h(\mu_i^*) \cdot \varepsilon_i \log \mu_i^* \cdot \omega_{ij,t} + O\Big(\omega_{ij,t}^2\Big), \tag{23}$$

where $h_{ij,t}^{D} \equiv (\mu_{ij,t}^{D} - 1)(\mu_{ij,t}^{D})^{-\varepsilon_i(1-\omega_{ij,t})}$ is the seller's profit at the distorted markup and elasticity. The gain is *first order* in $\omega_{ij,t}$: even a small amount of price-based learning allows the seller to charge a higher markup and earn strictly more than under disclosure. This is the profit the seller must forgo to sustain disclosure, and it determines whether disclosure is worth sustaining through relational enforcement.

5.3 Disclosure

The markup distortion creates a profit gain for the seller that can, in principle, be eliminated by voluntary information transmission—the kind of cost-justification communication that routinely accompanies price negotiations. We now ask: in which direction does disclosure generate enough seller profit gain to justify the cost of the information channel?

Upstream disclosure (seller → buyer) eliminates the markup distortion, recovering the first-order profit gain (23).

Downstream disclosure (buyer → seller) operates through a different channel. If the seller receives an additional signal about $q_{j,t}$ with precision $\kappa_{\text{add}} > 0$, its posterior variance falls by $\Delta\sigma^2_{\text{post}} \equiv \sigma^2_{\hat{q},\text{prior}} - \sigma^2_{\hat{q},\text{post}} > 0$, and the gain in the seller's expected profit is

$$\mathbf{E}\Big[\Pi^S_{k\ell,t}\Big(p^*(\hat{q}_j^+)\Big)\Big] - \mathbf{E}\Big[\Pi^S_{k\ell,t}\Big(p^*(\hat{q}_j)\Big)\Big] = \tfrac{1}{2}\,|h''(\mu_i^*)|\,\Delta\sigma^2_{\text{post}}\,\Gamma_{k\ell,t} + O\Big(\Big(\Delta\sigma^2_{\text{post}}\Big)^2\Big), \tag{24}$$

where $\hat{q}_j^+$ and $\hat{q}_j$ denote the seller's posterior with and without the additional signal. Because $h'(\mu_i^*) = 0$, the gain operates entirely through the second-order reduction in pricing error—an envelope-theorem effect. The buyer's surplus gain from the seller's improved pricing is also second order:

$$\mathbf{E}\Big[\Pi^B_{k\ell,t}\Big(p^*(\hat{q}_j^+)\Big)\Big] - \mathbf{E}\Big[\Pi^B_{k\ell,t}\Big(p^*(\hat{q}_j)\Big)\Big] = O\Big(\Delta\sigma^2_{\text{post}}\Big). \tag{25}$$

Proposition 1 (Asymmetric value of directional information)**.** *Consider a bilateral link between supplier ℓ in industry j and buyer k in industry i, and let $\omega_{ij,t} \in (0,1)$ denote the informational content of the bilateral price (equation* (19)*).*

(i) ***Upstream (seller → buyer).*** *The seller's profit gain from upstream disclosure is first order in $\omega_{ij,t}$: $h^D_{ij,t} - h(\mu^*_i) = h(\mu^*_i) \cdot \varepsilon_i \log \mu^*_i \cdot \omega_{ij,t} + O(\omega^2_{ij,t})$, with a strictly positive leading coefficient.*

(ii) ***Downstream (buyer → seller).*** *The gain in the seller's profit from an equivalent improvement in downstream information is $O(\omega^2_{ij,t})$—second order in the same perturbation parameter.*

(iii) ***No informative cheap talk.*** *Babbling is the unique equilibrium of any cheap-talk game over either party's belief.*

Proof. See Appendix C. □

The economic intuition follows from the two roles of the markup. On a non-disclosure link the seller charges $\mu^D_{ij,t} > \mu^*_i$, exploiting the buyer's lower perceived elasticity $\varepsilon_i(1 - \omega_{ij,t})$. Disclosing eliminates this exploitation: the seller loses $h^D_{ij,t} - h(\mu^*_i)$ today, which is first order in $\omega_{ij,t}$ because $\mu^D_{ij,t}$ is optimal at a different elasticity than μ^*_i. For the downstream direction, the seller already charges its privately optimal markup μ^*_i where $h'(\mu^*_i) = 0$; improving the seller's information about $q_{j,t}$ reduces pricing errors around this optimum by a second-order amount. This asymmetry—first order upstream, second order downstream—is what makes upstream disclosure worth sustaining through relational enforcement while downstream disclosure is not.[33]

Part (iii) closes the remaining route for voluntary downstream transmission. The buyer's surplus is strictly decreasing in p, and the seller optimally charges $p > \widehat{mc}_{j\ell,t}$ (since $\mu^*_i > 1$); for every realization of the buyer's type, the buyer prefers a lower price. This uniformly negative bias rules out any informative cheap-talk equilibrium by Crawford and Sobel (1982) (Theorem 1).[34] Cheap talk fails symmetrically in the upstream direction—the seller prefers the buyer to overestimate costs—but relational enforcement rescues upstream disclosure (Section 5.3). Could the same mechanism rescue downstream disclosure? In principle yes, but by part (ii) the seller's profit gain from downstream disclosure is second order in $\omega_{ij,t}$, too small to justify the same relational infrastructure. This completes the directional asymmetry: upstream disclosure generates a first-order profit gain that justifies relational enforcement; downstream disclosure does not.

We therefore model downstream belief transmission not as voluntary disclosure but as a non-strategic channel: sellers passively monitor buyers' observable activity and extract noisy non-price signals. Because this monitoring is non-strategic, its strength is naturally insensitive to relationship characteristics, consistent with the flat downstream heterogeneity pattern in the data (Table 4, Panel B). When we close the network belief system in Appendix F, these attention signals provide the baseline downstream information channel on all links.

[33] The asymmetry requires only $h'(\mu^*_i) = 0$ (private markup optimum) and $\mu^D_{ij,t} \neq \mu^*_i$ (distortion is non-zero). Both hold for any demand system with an interior markup optimum. If pricing is locally non-smooth or constrained, downstream information can generate first-order gains; we abstract from this because the downstream empirical pattern is a flat premium, captured through non-strategic monitoring.

[34] Network externalities do not break the argument: the buyer's bilateral payoff depends on the seller's price only through a monotone channel.

Disclosure cost, establishment condition, and threshold. Upstream disclosure requires the seller to truthfully communicate its posterior $\hat{q}_{j\ell,t}$. The disclosure decision is made period by period. Each period the seller chooses: disclose or not.

Payoffs. If the seller discloses and continues disclosing forever, its per-period profit is $h(\mu_i^*)\Gamma_{k\ell,t}$ and it pays a one-time setup cost $f \geq 0$ at establishment. If it deviates—stops disclosing this period—the buyer's demand is governed by the price signal, so the effective elasticity becomes $\varepsilon_i(1-\omega_{ij,t})$ and the seller optimally charges $\mu_{ij,t}^D$, earning $h_{ij,t}^D \equiv (\mu_{ij,t}^D - 1)(\mu_{ij,t}^D)^{-\varepsilon_i(1-\omega_{ij,t})}$ per unit of $\Gamma_{k\ell,t}$. Since $\mu_{ij,t}^D$ is optimal at elasticity $\varepsilon_i(1-\omega_{ij,t}) < \varepsilon_i$, we have $h_{ij,t}^D > h(\mu_i^*)$: deviating is profitable today. After detection, the relationship reverts permanently to non-disclosure, where the seller earns $h(\mu_{ij,t}^D)$ per period (at the true elasticity ε_i). Since $\mu_{ij,t}^D \neq \mu_i^*$ and h is maximized at μ_i^*, we have $h(\mu_{ij,t}^D) < h(\mu_i^*)$: non-disclosure is costly in all future periods.

IC. The seller discloses if and only if the value of disclosing forever exceeds the value of deviating today and reverting to non-disclosure:

$$(1 + R_{ij})\, h(\mu_i^*)\Gamma_{k\ell,t} - f \;\geq\; h_{ij,t}^D \Gamma_{k\ell,t} + R_{ij}\, h(\mu_{ij,t}^D)\Gamma_{k\ell,t},$$

where $R_{ij} \equiv \beta\rho_{ij}/(1-\beta\rho_{ij})$. Rearranging, disclosure is established if and only if:

$$f \;\leq\; [h(\mu_i^*) - h_{ij,t}^D]\Gamma_{k\ell,t} + R_{ij}\,[h(\mu_i^*) - h(\mu_{ij,t}^D)]\Gamma_{k\ell,t}. \tag{26}$$

The left-hand side is the setup cost. The right-hand side is the total seller profit gain from disclosing: $[h(\mu_i^*) - h_{ij,t}^D]\Gamma_{k\ell,t} < 0$ is the cost of not exploiting the distorted elasticity today, and $R_{ij}[h(\mu_i^*) - h(\mu_{ij,t}^D)]\Gamma_{k\ell,t} > 0$ is the PV of future profit gains. Disclosure requires the future gains to outweigh both the setup cost and today's foregone deviation profit.[35]

The setup cost f is uniform across links. Relationship durability ρ_{ij} is disciplined by the data: we compute it from average spell lengths and aggregate it to the industry-pair level. A higher ρ_{ij} raises R_{ij}, increasing the PV of future profit gains on the right-hand side and reducing the amortized setup cost f/R_{ij} on the left, making disclosure more likely on durable links. Identification of f is discussed in Section 6.1.

Proposition 2 (State-dependent transmission)**.** *For each link (k,ℓ) with $k \in i$ and $\ell \in j$, there exists a threshold $\bar{\sigma}_{q,j}^2$ such that disclosure is established if and only if $\sigma_{q,j}^2 \geq \bar{\sigma}_{q,j}^2$. The threshold is:*

(a) *decreasing in ρ_{ij} (relationship durability);*

(b) *decreasing in bilateral trade volume $\Gamma_{k\ell,t}$.*

The proof is in Appendix C. For part (a), a higher ρ_{ij} increases R_{ij}, raising the PV of future profit gains $R_{ij}\Delta\Pi_{k\ell,t}^S$ on the right-hand side of (26) and reducing the amortized setup cost f/R_{ij} on the left. Both effects make disclosure more likely on durable links.

Finally, the mechanism delivers the state dependence central to the empirics: higher $\sigma_{q,j}^2$ raises $\omega_{ij,t}$, which simultaneously increases $\Delta\Pi_{k\ell,t}^S$ (the future profit gain from disclosure) and $h_{ij,t}^D$ (today's deviation gain). The net effect on the RHS of (26) is positive to first order in ω, so disclosure becomes privately optimal precisely when aggregate uncertainty is high and information transmission is most valuable.

[35] The detection probability that sustains truthful reporting conditional on disclosure is subsumed in $\omega_{ij,t}$: higher detection reduces $\omega_{ij,t}$ and hence $h_{ij,t}^D$, making deviation less attractive. See Appendix C for the full derivation.

Remark 1. *A natural reduced-form extension is to allow the setup cost to depend on supplier breadth—for instance, replacing f with $f(1+b_j)^\gamma$, so that suppliers with broader networks face higher setup costs. This would add an additional prediction to the proposition: the disclosure threshold is increasing in b_j, generating a negative breadth gradient in belief premia alongside the empirically documented mirror pattern. We adopt the baseline specification $\gamma = 0$ (costs independent of breadth).*

Figure 4 summarizes the mechanism. As the price-signal weight ω rises, the distorted markup μ_D increases steeply—and with it, the seller's profit gain $\Delta\Pi^D$ from eliminating the distortion through disclosure. Disclosure is sustained whenever the present value of surplus $R_{ij} \cdot \Delta\Pi^D$ exceeds the setup cost f: links to the right of the threshold disclose, those to the left do not. The convexity of the markup distortion implies that the highest-ω links have the strongest incentive to disclose, endogenously resolving information frictions where they are most costly.

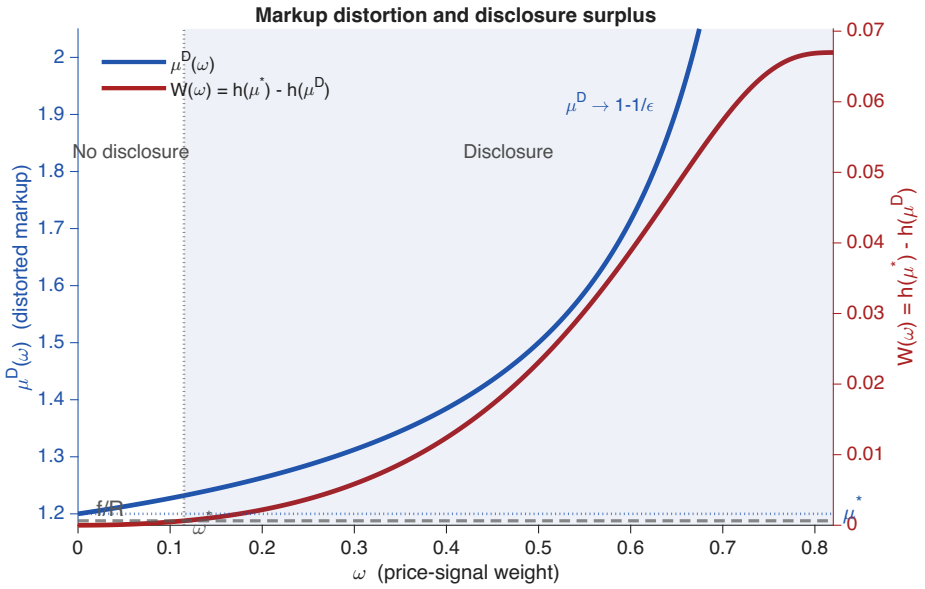


Figure 4: Markup distortion and disclosure surplus

Notes: Left axis (blue): distorted markup $\mu_D(\omega) = \varepsilon(1-\omega)/[\varepsilon(1-\omega)-1]$ as a function of the price-signal weight ω. The horizontal dotted line marks the efficient markup $\mu^* = \varepsilon/(\varepsilon-1) = 1.20$. Right axis (red): seller's net per-period gain from disclosure $\Delta\Pi^S(\omega) = h(\mu^*) - h(\mu^D)$, where $h(\mu) = (\mu-1)\mu^{-\varepsilon}$, net of the amortized deviation gain $[h^D - h(\mu^*)]/R_{ij}$. The vertical dashed line marks an illustrative disclosure threshold ω^*, above which the PV of net surplus $R_{ij} \cdot \Delta\Pi^S$ exceeds the setup cost f (horizontal dotted line, right axis). The shaded region indicates links that sustain disclosure. $\varepsilon = 6$.

6 Testing and quantifying the disclosure mechanism

This section quantifies the disclosure mechanism and its welfare implications. We proceed in three steps. First, we calibrate the model to the Italian production network and to the experimental moments documented in Sections 3–4. Second, we examine the equilibrium disclosure set and the model's fit to both targeted and non-targeted moments, including a validation against bilateral trade flows. Finally, we quantify the welfare cost of information frictions and evaluate counterfactual information regimes.

The section maintains a strict distinction between *calibration targets* and *validation moments*. The calibration uses a small set of moments that directly discipline the endogenous disclosure margin: the average upstream and downstream belief premia, the upstream price coefficient, the persistence gradient, and the size gradient. Other empirical patterns—including the bilateral trade elasticity, the sign reversal between partner-side and own-side gradients, and the activation of belief transmission during the 2021–2023 inflation episode—are treated as non-targeted validation checks. This distinction matters for interpreting the quantitative results: the model is disciplined by a compact set of moments tied directly to disclosure and then evaluated on additional dimensions that it was not forced to fit.

6.1 Calibration

We calibrate the model to the Italian production network and to the experimental moments documented above. The calibration strategy follows the empirical logic of the paper. The average upstream and downstream belief premia discipline the overall importance of direct communication. The persistence gradient disciplines how disclosure varies with relationship durability. These are the moments that directly identify the model's key endogenous margin: which links disclose and which rely on price-based inference.

Full details on moment construction, the objective function, and the computational algorithm are in Appendix S.

The production network is a 78×78 matrix A at the 2-digit ATECO level, constructed by pooling transaction flows over 2019–2023 and row-normalizing buyer industries' input shares. We set $\varepsilon = 6$, implying a steady-state markup of 20% following La'O and Tahbaz-Salehi (2022), and $\beta = 0.99$ at a quarterly frequency.

External parameters. Industry-level signal parameters are measured directly from the survey and price data. Table 7 summarizes their cross-industry distributions.

The variance of the common component is $\hat{\sigma}^2_{q,j} = \mathrm{Var}_t(\bar{g}^p_{j,t})$, the time-series variance of industry-average price growth. The idiosyncratic price-signal noise is

$$\hat{\sigma}^2_{\zeta,j} \;=\; \frac{1}{T}\sum_{t=1}^{T}\frac{1}{N_{jt}}\sum_{k\in j}\left(g^p_{k,t} - \bar{g}^p_{j,t}\right)^2,$$

the average within-industry, within-quarter cross-firm variance of price growth. The private-signal noise is identified from the industry-level price–cost pass-through:

$$\bar{g}^p_{j,t} \;=\; \mu_j \;+\; \alpha_j\,\bar{c}_{j,t} \;+\; e_{j,t}, \tag{27}$$

where $\bar{g}^p_{j,t}$ is industry-average realized price growth and $\bar{c}_{j,t}$ is industry-average realized input cost growth, both computed as cell means from the SIGE survey. The coefficient $\hat{\alpha}_j$ estimates the Kalman gain $\alpha_j = \sigma_q^2/(\sigma_q^2 + \sigma_\nu^2)$, from which we recover $\hat{\sigma}^2_{\nu,j} = \hat{\sigma}^2_{q,j}(1-\hat{\alpha}_j)/\hat{\alpha}_j$. This approach identifies the private-signal noise directly from the pricing equation, without relying on survey-based inflation expectations. The distinction matters because of the within-period timing described above: firms set prices from beginning-of-period information (their own signal about q_j), but report survey expectations after that quarter's transactions have occurred. The survey expectation therefore reflects the end-of-period posterior—which incorporates bilateral learning from supplier prices and disclosed messages—rather than the own-signal belief that

governs pricing. Using within-cell dispersion of survey expectations as a proxy for σ_ν^2 would understate the true private-signal noise, because bilateral learning causes beliefs to converge within industry cells. The price–cost pass-through, by contrast, is determined at the pricing stage before bilateral learning occurs, and therefore identifies α_j (and hence σ_ν^2) from the correct information set.

As a testable implication of own-signal pricing, the pass-through coefficient should not differ by treatment status, since the treatment affects the buyer's demand-side beliefs but not the seller's pricing signal. We find $\hat{\alpha}^{\text{ut}} = 0.52$ and $\hat{\alpha}^{\text{tr}} = 0.49$, with an interaction of -0.03 (s.e. 0.04)—indistinguishable from zero and consistent with own-signal pricing.

Link-level objects—expenditure shares, continuation probabilities ρ_{ij}, and relationship duration—are constructed from the administrative transaction data. For the treatment simulation, we scale the baseline $\hat{\sigma}^2_{\nu,j}$ by the ratio of treated-to-untreated within-cell forecast dispersion (median ratio 0.77), capturing the treatment's effect on the buyer's demand-side information without altering the pricing equation.

Table 7: External parameters

Parameter	Symbol	Median	P10	P90
Panel A: Signal variances (pp^2, 78 industries)				
Common-component variance	$\hat{\sigma}^2_{q,j}$	5.39	1.22	29.0
Idiosyncratic price noise	$\hat{\sigma}^2_{\zeta,j}$	20.1	6.70	43.2
Private-signal noise	$\hat{\sigma}^2_{\nu,j}$	5.94	1.90	10.3
Panel B: Price–cost pass-through (61 industries)				
Pass-through coefficient	$\hat{\alpha}_j$	0.48	0.20	0.71
Pass-through (pooled, ind. FE)	$\hat{\alpha}$	0.49 (s.e. 0.03)		
Treatment interaction	$\hat{\alpha}^{\text{tr}} - \hat{\alpha}^{\text{ut}}$	-0.03 (s.e. 0.04)		
Industries with $\Delta\tau > 0$		64 / 78		
Panel C: Link characteristics (2,794 active links)				
Expenditure share	a_{ij}	0.005	0.000	0.039
Continuation probability	ρ_{ij}	0.33	0.00	0.62
Price-signal weight	ω_{ij}	0.04	0.00	0.13

Notes: Panel A reports cross-industry percentiles of signal variance parameters. $\hat{\sigma}^2_{q,j}$ and $\hat{\sigma}^2_{\zeta,j}$ are measured from price data; $\hat{\sigma}^2_{\nu,j}$ is derived from the pass-through coefficient $\hat{\alpha}_j$ via $\hat{\sigma}^2_\nu = \hat{\sigma}^2_q(1-\hat{\alpha})/\hat{\alpha}$. Panel B reports the price–cost pass-through from regression (27). The treatment interaction tests the own-signal pricing prediction ($\hat{\alpha}^{\text{tr}} = \hat{\alpha}^{\text{ut}}$). Panel C reports link-level characteristics from the administrative transaction data. ω_{ij} is computed from the bilateral two-signal Kalman filter evaluated at equilibrium non-disclosure links. Fixed parameters: $\varepsilon = 6$, $\beta = 0.99$.

At median industry values, the price–cost pass-through is $\hat{\alpha} = 0.48$, implying that sellers pass roughly half of cost movements into prices. The implied private-signal noise is $\hat{\sigma}^2_\nu = 5.94$ pp^2 at the median, substantially larger than the common-component variance ($\hat{\sigma}^2_q = 5.39$ pp^2), confirming that the private signal is noisy. The aggregate price pass-through is $\theta'\alpha = 0.50$, constant across disclosure regimes since prices depend only on the seller's own signal. The median price-signal extraction weight across non-disclosure links is $\omega \approx 0.04$, with a right tail (P99 = 0.32): a handful of industries with high common-component variance relative to price noise generate links where prices carry meaningful information. The calibrated ratio $\hat{\sigma}^2_\zeta/\hat{\sigma}^2_q \approx 11.8$ at the median implies that idiosyncratic price noise substantially exceeds the common component, making the price signal relatively insensitive to changes in the supplier's precision—a property that limits the intensive-margin treatment effect on the price channel.

Disclosure costs. The model uses a uniform setup cost f across all links, with no per-period verification cost. The IC condition therefore reduces to

$$R_{ij}\, W(\omega_{p,j})\, \Gamma_{ij} \;\geq\; f, \tag{28}$$

where $W(\omega) = h(\mu^*) - h(\mu^D(\omega))$ is the cooperative surplus and $R_{ij} = \beta\rho_{ij}/(1-\beta\rho_{ij})$. This specification is motivated by the cross-sectional evidence: the duration gradient in the data is not monotone in the observed ρ_{ij} distribution (P10= 0.31, P90= 0.55), indicating that ρ_{ij} does not vary enough in the SIGE sample to identify a richer cost structure. With a uniform f, heterogeneity in disclosure is generated entirely by variation in R_{ij}, $\omega_{p,j}$, and Γ_{ij} across links.

Internally calibrated parameters. Two parameters are calibrated by minimizing a diagonal GMM criterion with weights equal to the inverse squared standard errors: the disclosure setup cost f and the price-noise scaling λ_ζ. All other signal parameters ($\hat{\sigma}^2_{q,j}$, $\hat{\sigma}^2_{\nu,j}$) are set externally. Table 8 summarizes the identification.

The mapping from parameters to moments is transparent. A higher f raises the IC threshold for disclosure, reducing the disclosure rate and hence the upstream belief premium; f is therefore identified by the upstream belief premium. A higher λ_ζ makes prices noisier, reducing the price-signal weight $\omega_{p,j}$ and hence the upstream price coefficient; λ_ζ is identified by the upstream price coefficient.

The two targeted moments are the upstream belief premium and the upstream price coefficient, giving an exactly identified system. Both moments are treatment effects on survey inflation expectations, which—as discussed above—reflect the end-of-period posterior that incorporates bilateral learning. The model counterparts are treatment-induced shifts in the demand-side belief weights ($\Delta\omega_{ij}$, $\Delta\psi_{ij}$), which is the correct mapping: the regressions measure how treated-partner signals enter the buyer's full posterior, and the model computes exactly this object through the bilateral Frisch–Waugh–Lovell decomposition. Note that these moments discipline the demand-side belief system (f and λ_ζ govern which links disclose and how informative prices are to buyers), while the pricing equation is identified separately from the pass-through regression (27). The two identification strategies target different stages of the within-period timing: the pass-through uses beginning-of-period pricing decisions; the targeted moments use end-of-period beliefs. The downstream belief premium, the persistence and size gradients, and the bilateral trade belief coefficient are reserved as non-targeted validation checks.

Table 8: Parameter identification

Parameter	Symbol	Moment	Source
Disclosure setup cost	f	Upstream belief premium	Table 3
Price-noise scaling	λ_ζ	Upstream price coeff.	Table 3

Targeted moments. Table 9 reports the fit to two targeted moments and four non-targeted validation moments. The objective function is a diagonal GMM criterion over the two targeted moments.

The calibrated parameters have natural interpretations. Rather than reporting f in raw GDP-normalised units, the economically meaningful object is the *cost-to-surplus ratio* $f\,/\,(R_{ij} \cdot W(\omega_{p,j}) \cdot \Gamma_{ij})$, which equals one at the disclosure threshold, exceeds one for non-disclosing links,

Table 9: Calibration results

Moment	Data	Model	S.E.
Panel A: Targeted moments (exactly identified)			
(1) Upstream belief premium ($T_j \times \pi^e_{j,t}$, up)	0.026	0.027	0.014
(2) Upstream price coeff. ($T_j \times g^p_{j,t}$)	0.010	0.011	0.007
Panel B: Non-targeted validation moments			
(3) Downstream belief premium ($T_j \times \pi^e_{j,t}$, dn)	0.079	0.036	0.021
(4) Duration gradient, belief (P90–P10)	0.133	0.043	0.023
(5) Size gradient, belief (P90–P10)	0.101	0.039	0.028
(6) Trade belief coeff. ($T_j \times \pi^e_{j,t}$, trade, lagged)	0.063	0.032	0.022

Notes: Panel A reports the two moments targeted in calibration; the system is exactly identified. Panel B reports non-targeted validation moments. Moments (3)–(5) are discussed in Section 4; moment (6) is discussed in Section 6.2. Calibrated parameters: $\lambda_\zeta = 3.16$ ($\hat{\sigma}^2_\zeta$ scaled from 20.1 to 63.4 pp^2); f normalised to GDP units. External parameters (medians): $\hat{\sigma}^2_\nu = 5.94$ (from pass-through $\hat{\alpha} = 0.48$), $\hat{\sigma}^2_\zeta = 20.1$, $\hat{\sigma}^2_q = 5.39$. Fixed: $\varepsilon = 6$, $\beta = 0.99$, $\gamma = 0$.

and falls below one for disclosing links. At the median non-disclosing link the setup cost substantially exceeds the present value of cooperative surplus, placing the link well inside the non-disclosure region. At the median disclosing link the cooperative surplus exceeds the setup cost, making disclosure privately optimal. The sharp contrast between these two groups reflects the strong gradient in R_{ij} and Γ_{ij} across the link distribution, and is precisely what generates the steep disclosure gradient by expenditure share in Figure 5.

An alternative way to gauge the magnitude of f is relative to the present discounted value of bilateral profits, $(1+R_{ij})\, h(\mu^*)\, \Gamma_{ij}$. At the median disclosing link, f is less than one percent of the PDV of bilateral profits. These are small numbers in absolute terms, but the cooperative surplus on which the IC condition binds—the markup gain from eliminating price-based inference—is itself a small fraction of total profits. Even a sub-percent setup cost therefore exceeds the disclosure gain on links where prices are relatively uninformative (small $\omega_{p,j}$) or trade volumes are low (small a_{ij}).

$\lambda_\zeta = 3.16$ implies that the effective idiosyncratic price noise is roughly $3\times$ larger than the raw within-industry price dispersion. This scaling is economically plausible: the administrative price data reflects demand shocks, product-mix variation, and measurement error that are orthogonal to the common component in the model. The corrected ratio $\hat{\sigma}^2_\zeta/\hat{\sigma}^2_q \approx 11.8$ at the median implies that prices carry limited information about the common component, making the price signal relatively insensitive to changes in supplier precision—a property that limits the intensive-margin treatment effect on the price channel and is consistent with the small and imprecise upstream price coefficient in Table 3.

6.2 Equilibrium disclosure and model fit

A model with two calibrated parameters and a single disclosure margin accounts for roughly one-third to one-half of the observed cross-link heterogeneity in four out-of-sample validation moments it was not asked to fit. Table 9 reports the full picture.

The bilateral trade belief coefficient (moment 6) is the cleanest out-of-sample test: it uses a different outcome variable (log trade value rather than beliefs), a different sample (administrative transaction data rather than survey responses), and a different identification strategy (cross-supplier within buyer). The model predicts 0.032 against a data value of 0.063, within 1.4

standard errors—confirming that the disclosure mechanism has real allocative consequences of the right order of magnitude.

The size gradient (moment 5) is captured at 39% (0.039 vs. 0.101, P90–P10). The model generates this gradient through the cooperative surplus $W(\omega_j) \cdot \Gamma_{ij}$, which rises with expenditure share a_{ij}: larger links have more to gain from disclosure and more readily cross the IC threshold. The downstream belief premium (moment 3) is captured at 46% (0.036 vs. 0.079): in the model, downstream transmission is frictionless cheap-talk, so the premium equals the treatment-induced precision gain weighted by the Kalman gain.

The duration gradient (moment 4) is the hardest moment at 32% (0.043 vs. 0.133). The IC condition amplifies the cooperative surplus on persistent links through $R_{ij} = \beta\rho_{ij}/(1-\beta\rho_{ij})$, but the observed ρ_{ij} distribution in the SIGE sample is compressed (P10= 0.31, P90= 0.55), limiting the gradient the model can generate to at most 0.097. The gap between 0.097 and 0.133 reflects relationship-specific communication quality—the tendency for durable partnerships to develop richer, lower-cost information channels—that a single-margin model cannot capture. This points naturally to the next step: a model where f itself falls with relationship age, generating steeper duration gradients without altering the welfare mechanism.

The pattern across validation moments is informative in its own right. The model does well on the margins most directly tied to the IC condition—bilateral trade and link size—and less well on the duration margin, where the compressed ρ distribution limits its reach. This ordering is exactly what the theory predicts: the single margin of disclosure versus non-disclosure captures the *level* of heterogeneity across link types, but not the *within-type* variation driven by communication quality. We interpret the baseline as providing a disciplined lower bound on the welfare consequences of imperfect disclosure: a richer model would shift more links into disclosure and generate larger welfare gains, not smaller ones.

Equilibrium disclosure. The calibrated model implies that disclosure is substantial but far from universal. Table 10 shows that approximately 23% of active links sustain disclosure, with a strong gradient by link size: 6% in the bottom quartile of expenditure share versus 52% in the top quartile. Disclosure also responds substantially to aggregate uncertainty, rising from 12% to 41% when σ_q^2 doubles. This strong state dependence is consistent with the empirical evidence in Table 6, where belief premia increase substantially during high-inflation quarters, and reflects the fact that many links sit near the disclosure threshold.

Table 10: Equilibrium disclosure properties

	Disclosure rate
All active links	23.2%
Bottom quartile (expenditure share)	6%
Top quartile (expenditure share)	52%
Low uncertainty ($0.5 \times \bar{\sigma}_q^2$)	12%
High uncertainty ($2 \times \bar{\sigma}_q^2$)	41%

Figure 5 reports the equilibrium disclosure rate by quartile of expenditure share a_{ij}. Disclosure rises steeply with link size: 6% in the bottom quartile versus 52% in the top quartile, consistent with the larger cooperative surplus $W(\omega_j) \cdot \Gamma_{ij}$ on economically important links. Appendix R reports the corresponding gradient by relationship persistence, which is non-monotone: disclosure rises from near zero at low ρ to above 40% around the median, then declines in the

right tail because high-persistence links tend to have below-average expenditure shares, so the IC condition—which binds on the product $R_{ij} \cdot W(\omega_j) \cdot \Gamma_{ij}$—is not monotone in ρ_{ij} alone.

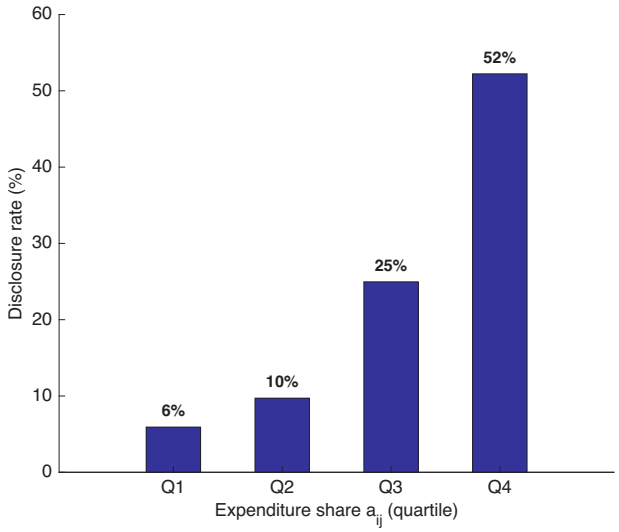


Figure 5: Equilibrium disclosure rate by expenditure share

Notes: Each bar reports the disclosure rate among active links in the calibrated model, by quartile of expenditure share a_{ij}. Calibrated parameters as in Table 9.

Figure 6, Panel A, illustrates the disclosure equilibrium. The gray line plots the cumulative distribution of the price-signal weight ω_{ij} across all 6,011 active links, equivalent to the no-disclosure counterfactual. The black line plots the distribution over the 4,616 non-disclosure links that remain in equilibrium. The gap between the two CDFs represents links that choose to disclose: these are disproportionately high-ω links, where the markup distortion from price-based learning is largest. Disclosure therefore resolves information frictions where they are most costly, shifting the residual ω distribution to the left.

Bilateral trade and the disclosure channel

A central implication of the model is that information disclosure has real effects on firm-to-firm allocations. When a supplier is better informed about aggregate conditions, the buyer can use the supplier's communicated beliefs as a credible signal about cost fundamentals, reducing the uncertainty that would otherwise depress bilateral trade. We test this implication using administrative bilateral transaction data, which provide a real-allocation margin that is distinct from the survey-based expectations evidence.

Estimating equation. We estimate a dyad-level specification with log bilateral trade value on the left-hand side, restricting to untreated buyers ($T_i = 0$):

$$\begin{aligned}\log \Gamma_{k\ell,t} = \alpha_k + \delta_{s(j),t} + \alpha_{s(k),t} + \alpha_{\text{size}(k),t} + \alpha_{\text{area}(k),t} + \alpha_{T_j,t} \\ + \beta_0^p \, g_{j,t}^p + \beta_1^p \, (T_j \times g_{j,t}^p) + \beta_0^e \, \pi_{j,t}^e + \beta_1^e \, (T_j \times \pi_{j,t}^e) + \varepsilon_{k\ell,t},\end{aligned} \tag{29}$$

where $\Gamma_{k\ell,t}$ is the transaction value between buyer k and supplier ℓ in quarter t, $g_{j,t}^p$ is the supplier's price growth, and $\pi_{j,t}^e$ is the supplier's 12-month-ahead inflation expectation. The fixed effects include buyer fixed effects α_k; supplier-industry×quarter fixed effects $\delta_{s(j),t}$; and

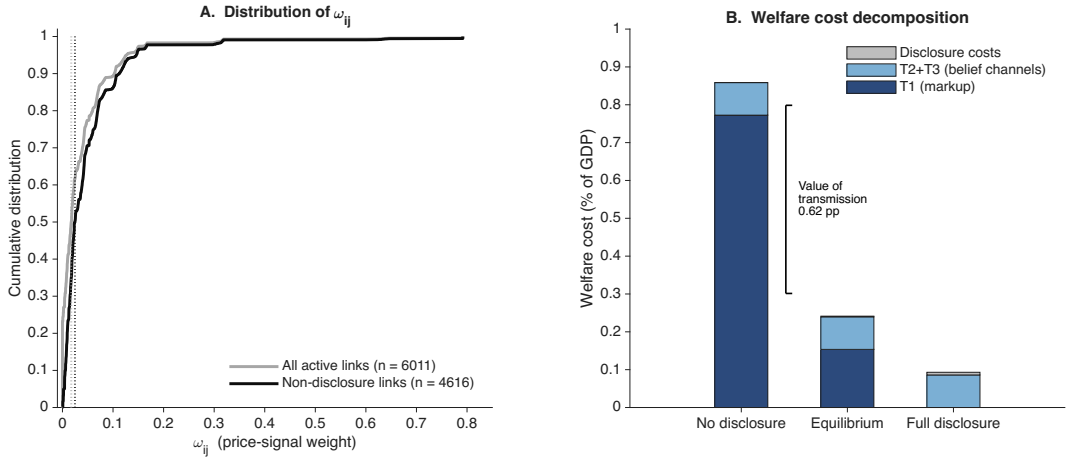


Figure 6: Disclosure equilibrium and welfare

Notes: Panel A plots the CDF of the price-signal weight ω_{ij} across active links. Gray: all 6,011 active links (no-disclosure counterfactual). Black: the 4,616 non-disclosure links in the equilibrium disclosure set. The gap represents links that choose to disclose—disproportionately high-ω links where the markup distortion is largest. Panel B reports the consumption-equivalent welfare cost of information frictions decomposed into the first-order markup channel (T1, dark) and the second-order information channels (T2+T3, light), across three regimes. The bracket indicates the value of existing information transmission.

buyer sector×quarter, size×quarter, and area×quarter fixed effects. The term $\alpha_{T_j,t}$ denotes T_j×quarter fixed effects.

The specification includes both the supplier's price growth and inflation expectations, interacted with treatment status, to distinguish the price channel from the belief channel. The supplier-industry×quarter fixed effects absorb industry-specific trends in trade volumes that might otherwise confound the treatment effect. The identifying variation is cross-supplier within a buyer: conditional on buyer k's overall procurement in quarter t and the supplier's industry trend, does k trade differentially with treated versus untreated suppliers as a function of the supplier's price growth and reported expectations?

Model predictions. The CES demand structure delivers predictions for both coefficients of interest (Appendix R.2). The model predicts $\beta_1^p > 0$: on non-disclosure links, the treatment makes supplier prices more informative about fundamentals, dampening the buyer's substitution response. It also predicts $\beta_1^e > 0$: the treatment raises the precision of the supplier's communicated beliefs, increasing the weight the buyer places on this signal when evaluating the outside option. Both predictions arise from the same mechanism—treatment improves the quality of supplier-generated signals—but they operate on different margins of the buyer's information set.

Results. Table 11 reports the results. The key finding is a sharp asymmetry between the price and belief channels. The treated-supplier price interaction $\hat{\beta}_1^p$ is small and statistically insignificant across all specifications: conditional on supplier-industry×quarter fixed effects, the treatment does not differentially affect the trade-value response to supplier prices. By contrast, the treated-supplier belief interaction $\hat{\beta}_1^e$ is large and highly significant.

Under uniform weights (column 1), $\hat{\beta}_1^e = 0.104$ (s.e. 0.031): a one-percentage-point increase in a treated supplier's inflation expectation raises bilateral trade by approximately 10 percent more than the same increase from an untreated supplier. The effect is stable across all four weighting

schemes (columns 1–4), with t-statistics ranging from 2.2 to 3.5. The base effect of supplier expectations is negative ($\hat{\beta}_0^e \approx -0.10$): buyers trade less with suppliers who expect higher inflation, consistent with the model's prediction that uncertainty about supplier-side fundamentals depresses bilateral trade. The treatment interaction almost exactly offsets this penalty, so that for treated suppliers the net effect of higher expectations on trade is approximately zero ($\hat{\beta}_0^e + \hat{\beta}_1^e \approx 0$).

This pattern has a natural interpretation within the model. Higher supplier inflation expectations signal higher expected costs, which would normally reduce a buyer's willingness to trade. But when the supplier is treated—and therefore better informed—the buyer interprets the supplier's expectations as a credible signal rather than noise. The treatment makes the supplier's communication informative, sustaining trade that would otherwise be reduced under uncertainty. This is the disclosure mechanism operating through beliefs rather than prices.

Robustness. The belief coefficient is robust to a comprehensive set of specification checks (Appendix R.3). It survives trimming small transactions across a range of thresholds, with point estimates between 0.08 and 0.11 and t-statistics consistently above 2.9. It strengthens on established links lasting four or more quarters, indicating that the disclosure channel operates most strongly in relationships with meaningful transaction volumes. Under alternative clustering schemes the coefficient remains significant at the 1% level.

To assess how much of the effect reflects persistent cross-link differences, we condition on lagged bilateral trade value $\log \Gamma_{k\ell,t-1}$. The belief coefficient attenuates to $\hat{\beta}_1^e = 0.063$ ($t = 2.86$): roughly two-thirds of the baseline estimate is absorbed by the persistent component of trade, but the incremental effect remains significant. Under dyad fixed effects, which absorb all time-invariant link heterogeneity, the coefficient is insignificant ($\hat{\beta}_1^e = 0.003$), consistent with the disclosure mechanism operating through persistent cross-link differences—as the model predicts, since treatment status is fixed and disclosure is a level decision—rather than through quarter-to-quarter fluctuations in survey responses.

The absence of a price channel in the bilateral trade data—despite its presence in the expectations regressions at marginal significance—is informative. In the model, the price channel operates through the buyer's Bayesian updating of beliefs about $Q_{j,t}$ from the bilateral price signal. The bilateral trade regression conditions on supplier-industry×quarter fixed effects, which absorb industry-level price variation and leave only within-industry cross-supplier price differences. The belief channel, by contrast, operates through the supplier's disclosed expectations, which vary at the supplier level and survive industry controls.

We therefore treat the bilateral trade evidence as complementary to the expectations-formation results: it confirms that information transmission has real allocative consequences and that these consequences operate primarily through the belief channel. The model-implied counterpart to the lagged belief coefficient is reported as the non-targeted validation moment in Panel B of Table 9.

6.3 Welfare and counterfactuals

We now turn to the quantitative stakes of the mechanism. The central welfare implication of the model is that price-based inference generates a first-order markup distortion on non-disclosure links. Disclosure removes this distortion link by link. The welfare question is therefore not primarily whether firms forecast aggregate inflation more accurately, but how much of the

Table 11: Bilateral trade and the disclosure channel

	Dep. var.: $\log \Gamma_{k\ell,t}$			
	(1) Uniform	(2) Breadth	(3) Value	(4) Share
$g^p_{j,t}$	0.021**	0.022*	−0.001	0.006
	(0.010)	(0.013)	(0.010)	(0.009)
$T_j \times g^p_{j,t}$	0.016	0.018	0.017	0.023**
	(0.012)	(0.015)	(0.011)	(0.011)
$\pi^e_{j,t}$	−0.109***	−0.144***	−0.108***	−0.100***
	(0.027)	(0.035)	(0.031)	(0.023)
$T_j \times \pi^e_{j,t}$	0.104***	0.110***	0.086**	0.103***
	(0.031)	(0.041)	(0.039)	(0.029)
Sup. industry × qtr FE	Yes	Yes	Yes	Yes
Observations	54,912	54,912	54,912	54,912
R^2	0.604	0.663	0.809	0.763

Notes: Dependent variable is $\log \Gamma_{k\ell,t}$ (log bilateral trade value, winsorized at the 1st and 99th percentiles). Sample restricted to untreated buyers ($T_i = 0$), upstream links only. All specifications include buyer fixed effects, buyer sector×quarter, size×quarter, and area×quarter fixed effects, T_j×quarter fixed effects, and supplier-industry×quarter fixed effects. Columns differ in link weights: uniform (1), supplier breadth (2), transaction value (3), supplier sales share (4). Standard errors (in parentheses) clustered at the buyer firm level. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

network remains in the non-disclosure regime where prices must carry information.

Welfare decomposition. Panel B of Figure 6 and Table 12 report the consumption-equivalent welfare cost of information frictions, decomposed into a first-order *markup channel* (T1), an aggregate *belief error channel* (T2), and a within-industry *dispersion channel* (T3), following La'O and Tahbaz-Salehi (2022).[36]

The equilibrium cost is 0.24% of GDP. The majority of this (0.15 pp) arises from the markup channel: non-disclosure links sustain distorted markups that generate first-order welfare losses. The aggregate belief-error channel contributes 0.026% and within-industry dispersion 0.060%—together roughly one-third of the total. Crucially, T2 and T3 are constant across disclosure regimes: they reflect the pricing errors inherent in own-signal pricing and do not respond to changes in the information regime. Only T1 and disclosure costs vary with the extent of disclosure. The markup cost T1 is moderately sensitive to the right tail of the ω distribution: capping the markup distortion at $\omega \leq 0.50$ reduces T1 by roughly 4%, confirming that the welfare estimate is not driven by extreme links.

Removing all belief channels raises the total cost to 0.86% of GDP. The *value of existing information transmission*—the net welfare gain from equilibrium disclosure relative to no belief channels, inclusive of disclosure costs—is 0.62 percentage points of GDP, overwhelmingly through the markup channel: disclosure eliminates the markup distortion on the 23% of links that sustain it, reducing T1 from 0.77% to 0.15%.

[36] All channels are expressed as fractions of GDP using Domar weights. T1 captures the first-order surplus loss from distorted markups on non-disclosure links; T2 captures the loss from aggregate belief errors; T3 captures within-industry dispersion. A Monte Carlo verification with 100,000 replications confirms that the analytical welfare formula matches simulated welfare to within 4% accuracy; the residual reflects higher-order terms in the Taylor expansion.

Table 12: Welfare cost of information frictions (% of GDP)

	T1 (markup)	T2 (agg. error)	T3 (dispersion)	Disc. costs	Total
Panel A: Information regimes					
No belief channels	0.77	0.03	0.06	0.00	0.86
Equilibrium disclosure	0.15	0.03	0.06	0.00	0.24
Full disclosure	0.00	0.03	0.06	0.01	0.09
Value of transmission	0.62	0.00	0.00	−0.00	0.62
Panel B: Treatment counterfactual					
SIGE treatment	0.16	0.02	0.06	0.00	0.24
Treatment effect	+0.01	−0.01	−0.01	+0.00	−0.00
Links switched			141 / 6,011		
Panel C: Aggregate stabilization counterfactual					
Industry-specific only ($\mathrm{Var}(\bar{q}) = 0$)	0.15	0.02	0.04	0.00	0.21
Due to aggregate component	0.00	0.00	0.02	0.00	0.03

Notes: All entries are consumption-equivalent welfare costs expressed as percentages of GDP, computed using Domar weights. Panel A compares information regimes: "No belief channels" sets $\Psi^R = \Phi = 0$; "Full disclosure" forces all links to disclose; "Value of transmission" is the net gain from equilibrium disclosure relative to no belief channels, inclusive of disclosure costs. Panel B evaluates the SIGE experiment. Panel C decomposes each industry's cost shock q_j into a consumption-weighted aggregate component $\bar{q} = \sum_i \theta_i q_i$ and an industry-specific residual. "Industry-specific only" re-solves the model with $\mathrm{Var}(\bar{q}) = 0$, retaining only residual industry-level uncertainty.

Price pass-through and the distinction from La'O and Tahbaz-Salehi. A key feature of the welfare decomposition is that only T1 varies with the disclosure regime, while T2 and T3 are *constant* across all information regimes. This follows from the own-signal pricing structure: the seller's price depends only on its private signal, not on the network of bilateral learning. Formally, the pricing matrix is $M^{\text{price}} = \mathrm{diag}(\alpha)$, where $\alpha_j = \sigma_q^2/(\sigma_q^2 + \sigma_\nu^2)$, and the aggregate price pass-through is $\theta'\alpha = 0.50$—*invariant to the disclosure regime.* Network learning (price signals, disclosed beliefs, the Leontief inverse $\mathcal{H}$) governs the buyer's demand and the bilateral surplus, but not the seller's price.

This structure creates a sharp separation between the welfare cost of information frictions in our model and in La'O and Tahbaz-Salehi (2022) (LTS). In LTS, prices are set by Bayesian agents who combine all available signals; the welfare cost arises because dispersed information generates aggregate output-gap volatility (our T2) and within-industry misallocation (our T3). These terms are second order in belief errors and vanish as information improves. In our model, the dominant welfare cost is T1—a first-order markup distortion that arises because the buyer's demand depends on the *difference* between disclosed and price-inferred beliefs. T1 is a feature of the pricing game, not the belief-formation problem: no amount of better buyer attention eliminates it, making disclosure's welfare effects first order rather than second order.

Quantitatively, the welfare decomposition in Table 12 confirms this distinction. The LTS-type terms (T2+T3) sum to 0.09% of GDP and are unchanged across regimes. The markup wedge T1—driven by buyer–seller information asymmetry—ranges from 0.77% (no disclosure) to 0.00% (full disclosure), an order of magnitude larger. The value of existing information transmission (0.62 pp) comes entirely from T1. A model that treated information frictions as generating only second-order welfare costs would miss the dominant channel.

Treatment counterfactual and the markup channel. The SIGE treatment—which reduces treated firms' private-signal noise by a factor of 0.77 (the treated-to-untreated dispersion ratio)—yields a small welfare improvement (-0.004 pp, i.e., slightly lower costs). The effect is economically negligible, reflecting the interaction between the intensive and extensive margins of the price channel.

Treatment improves firms' private precision (α_j rises), which has two opposing effects. On the *extensive margin*, better precision raises the cooperative surplus from disclosure, potentially switching links into the disclosure set—a welfare gain, as each switching link eliminates its markup distortion. On the *intensive margin*, higher precision makes the supplier's price more informative about fundamentals, raising $\omega_{ij,t}$ on every non-disclosure link—a welfare loss through higher markup distortions.

In the calibrated economy, the extensive margin is modest: 141 of 6,011 active links switch into disclosure. The intensive margin partially offsets this: on the remaining non-disclosure links, ω rises from 0.050 to 0.055, generating a small increase in markup distortions. The belief channels (T2, T3) improve slightly under treatment (-0.008 and -0.005 pp respectively), roughly offsetting the T1 increase, so the net welfare effect is approximately zero.

Two implications emerge. First, the welfare benefits of information provision depend not on the *level* of signal quality but on whether quality improvements cross the disclosure threshold—a prediction unique to models with endogenous information sharing. Marginal improvements in firm-level precision have negligible aggregate welfare effects unless they trigger new disclosure relationships. Second, policies that reduce disclosure costs for specific links—for instance, standardized reporting frameworks that lower setup costs for established relationships—may be more effective than broad-based information provision at generating welfare gains through the network.

Aggregate versus industry-specific cost uncertainty. The markup distortion on non-disclosure links arises from buyer uncertainty about the common component of the supplier's costs, $q_{j,t}$, which reflects both an aggregate component shared across industries and an industry-specific residual. How much of the equilibrium welfare cost could be eliminated by perfectly stabilizing the aggregate price level?

Panel C of Table 12 reports the result. We decompose each industry's common component q_j into a consumption-weighted aggregate $\bar{q} = \sum_i \theta_i q_i$ and an orthogonal industry-specific residual, then re-solve the model with $\text{Var}(\bar{q}) = 0$. Perfect aggregate stabilization reduces the equilibrium welfare cost from 0.24% to 0.21% of GDP—a modest reduction of 0.03 pp, accounting for 12% of the total cost.

The small net effect masks two large opposing forces, which we isolate using a fixed-R counterfactual. The *direct effect* of removing aggregate uncertainty—holding the disclosure set fixed—reduces the welfare cost by 0.09 pp, as ω_{ij} falls on all non-disclosure links. But the *disclosure margin response* offsets two-thirds of this gain: with less aggregate uncertainty, the cooperative surplus from disclosure shrinks, and the disclosure rate falls from 23% to 12%. The newly non-disclosing links revert to price-based inference, generating markup distortions that erode the direct benefit. The net effect (0.03 pp) is the small residual of these two forces.

This offsetting mechanism is a distinctive prediction of the model. It arises because disclosure is endogenous: reducing cost uncertainty lowers the returns to disclosure, causing some links to exit the disclosure set. This "crowding out" effect is absent in models with exogenous information structures, where reducing uncertainty unambiguously improves welfare.

This contrasts sharply with the no-disclosure benchmark: in a world without any belief channels, aggregate uncertainty accounts for 0.41 pp of the 0.86% no-disclosure cost (48%). Without disclosure, there is no margin to crowd out: prices are the only information channel, and removing the aggregate component of σ_q^2 directly reduces ω_{ij} on every link. Disclosure effectively insulates welfare from the aggregate component by concentrating non-disclosure on links where industry-specific variation dominates—and by creating an endogenous margin that responds adversely to aggregate stabilization.

Comparison with endogenous attention. The crowding-out of disclosure by aggregate stabilization resembles, at a high level, the reallocation of attention in rational-inattention models (Afrouzi, 2024): in both environments, reducing aggregate uncertainty weakens an endogenous information margin and thereby offsets part of the direct welfare gain. The mechanism here is different in two respects. First, the offset operates through the first-order markup wedge (T1), rather than through second-order pricing errors (T2, T3). Second, the disclosure margin is governed by a private IC constraint rather than by efficient reallocation of a scarce cognitive resource. Firms exit disclosure when bilateral surplus falls below the setup cost, even though disclosure may remain socially valuable because it reduces downstream markup distortions. The crowding-out induced by stabilization is therefore socially inefficient in our environment, which helps explain why the residual welfare gain from aggregate stabilization is small.

7 Conclusions

This paper studies how firms learn about aggregate inflation from their trading partners. Using a randomized information treatment that creates persistent differences in information precision, merged with the universe of business-to-business transactions in Italy, we provide causal evidence that firms learn from trading partners through both beliefs and prices, and that the relative importance of these channels varies systematically across relationships. Belief-based learning is strongest in durable and economically important supplier relationships, while price-based learning is stronger for suppliers with broader networks. These gradients are fundamentally supplier-side rather than buyer-side, pointing to the partner's incentives to communicate as the key margin.

To interpret these patterns, we develop a model of strategic belief disclosure in production networks. Suppliers choose whether to communicate information directly to buyers or leave them to infer it from prices. Disclosure is easier to sustain in durable, high-surplus relationships and harder to sustain when suppliers manage broader networks, generating the heterogeneity observed in the data. The model also implies that the belief channel, rather than the price channel, should matter for bilateral trade volumes. We confirm this prediction in the data: better-informed suppliers' beliefs predict bilateral trade, whereas suppliers' prices do not.

Embedding the bilateral mechanism in the full production network, we calibrate the model to two experimental moments—the upstream belief premium and the upstream price coefficient—in an exactly-identified GMM system. Existing information transmission recovers 0.62 percentage points of GDP relative to a world without belief channels. Most of this gain operates through the markup channel: disclosure removes the price-inference wedge on non-disclosure links and thereby improves bilateral allocations. The residual equilibrium welfare cost is 0.24% of GDP, with second-order pricing errors accounting for a smaller share of the total. Aggregate price stabilization delivers only a modest additional gain, because reducing aggregate cost volatility

induces firms to exit disclosure relationships and thereby leaves the dominant bilateral distortion largely intact.

Our findings point to a broader lesson. In production networks, the key information friction is not only how well firms estimate aggregate conditions, but also whether trading partners communicate directly or instead rely on price inference. That distinction matters for welfare because price inference distorts markups link by link, and it matters for policy because aggregate stabilization has limited traction on a bilateral communication problem. More broadly, the analysis isolates one particular margin of learning in production networks: endogenous disclosure within a static relational environment. An important next step is to understand how this margin interacts with endogenous attention and with the persistence of buyer–seller relationships over time, both of which may shape the credibility and value of information transmission. It would also be useful to assess how far the mechanism travels across countries and institutional settings, where differences in network structure, inflation experience, and communication norms may change the relative importance of direct beliefs and price inference.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Acemoglu, D. and A. Tahbaz-Salehi (2025). The macroeconomics of supply chain disruptions. *Review of Economic Studies 92*(2), 656–695.

Afrouzi, H. (2024). Strategic inattention, inflation dynamics, and the nonneutrality of money. *Journal of Political Economy 132*(10), 3378–3420.

Amador, M. and P.-O. Weill (2010). Learning from prices: Public communication and welfare. *Journal of Political Economy 118*(5), 866–907.

Amador, M. and P.-O. Weill (2012). Learning from private and public observations of others' actions. *Journal of Economic Theory 147*(3), 910–940.

Andrade, P., O. Coibion, E. Gautier, and Y. Gorodnichenko (2022). No firm is an island? how industry conditions shape firms' expectations. *Journal of Monetary Economics 125*, 40–56.

Angeletos, G.-M. and Z. Huo (2021). Myopia and anchoring. *American Economic Review 111*(4), 1166–1200.

Angeletos, G.-M. and J. La' O (2020). Optimal monetary policy with informational frictions. *Journal of Political Economy 128*(3), 1027–1064.

Angeletos, G.-M. and C. Lian (2023). The information structure of the economy. *American Economic Review 113*(9), 2364–2411.

Babus, A. and P. Kondor (2018). Trading and information diffusion in over-the-counter markets. *Econometrica 86*(5), 1727–1769.

Bui, H., Z. Huo, A. A. Levchenko, and N. Pandalai-Nayar (2024). Noisy global value chains. Mimeo, Yale University.

Carvalho, V. M. and A. Tahbaz-Salehi (2019). Production networks: A primer. *Annual Review of Economics 11*, 635–663.

Coibion, O. and Y. Gorodnichenko (2025). The new causal macroeconomics of surveys and experiments. *Journal of the European Economic Association*. forthcoming.

Coibion, O., Y. Gorodnichenko, and T. Ropele (2020). Inflation expectations and firm decisions: New causal evidence. *The Quarterly Journal of Economics 135*(1), 165–219.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–1451.

Driscoll, J. C. and A. C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics 80*(4), 549–560.

Fang, Y., Z. Huo, M. Li, and J. Wu (2025). Inattentive network: Theory and evidence from EDGAR. Mimeo, Yale University.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics 129*(4), 1661–1710.

Gaballo, G. and L. Paciello (2025). Spending allocation under nominal uncertainty: A model of effective price rigidity. *Journal of the European Economic Association 23*(6), 2376–2413.

Ghassibe, M. and A. Nakov (2025). Business cycles with pricing cascades. CEPR Discussion Paper 20605, CEPR.

Hajdini, I., S. Kumar, S. Malik, J. J. Norris, and M. Pedemonte (2025). Supply chain networks and the macroeconomic expectations of firms. Mimeo, Federal Reserve Bank of Cleveland.

Huo, Z., M. Li, J. Wu, and H. Xiong (2025). Micro shocks and macro fluctuations in the information network. Mimeo, Yale University.

Janssen, M. and S. Shelegia (2015). Consumer Search and Double Marginalization. *American Economic Review 105*(6), 1683–1710.

La'O, J. and A. Tahbaz-Salehi (2022). Optimal monetary policy in production networks. *Econometrica 90*(3), 1295–1336.

Lorenzoni, G. (2009). A theory of demand shocks. *American Economic Review 99*(5), 2050–2084.

Lorenzoni, G. (2010). Optimal monetary policy with uncertain fundamentals and dispersed information. *Review of Economic Studies 77*(1), 305–338.

Lucas, R. E. (1972). Expectations and the neutrality of money. *Journal of Economic Theory 4*(2), 103–124.

Mankiw, N. G. and R. Reis (2002). Sticky information versus sticky prices: A proposal to replace the new Keynesian phillips curve. *Quarterly Journal of Economics 117*(4), 1295–1328.

Morris, S. and H. S. Shin (2002). Social value of public information. *American Economic Review 92*(5), 1521–1534.

Nikolakoudis, A. (2024). Information diffusion in networks: A theory of learning from prices. Mimeo, University of Chicago.

Rosolia, A. (2024). Do firms act on their inflation expectations? another look at italian firms. *Journal of Political Economy Macroeconomics 2*(4), 651–686.

Weber, M., B. Candia, H. Afrouzi, T. Ropele, R. Lluberas, S. Frache, B. Meyer, S. Kumar, Y. Gorodnichenko, D. Georgarakos, et al. (2025). Tell me something i don't already know: Learning in low-and high-inflation settings. *Econometrica 93*(1), 229–264.

A From model to regression

This appendix derives the regression specification (6) from the model's belief formation structure.

Posterior and signal weights. By equation (51), firm k in industry i forms a posterior about the common component $q_{i,t}$ as a weighted combination of all available signals:

$$\hat{q}_{ik,t} \;=\; \kappa_i \, s_{q,k,t} \;+\; \sum_{j \in \mathcal{N}_i^{\mathrm{up}}} w_{ij,t} \, \tilde{s}_{ij,k,t} \;+\; (\text{downstream and attention signals}), \tag{30}$$

where $\tilde{s}_{ij,k,t}$ denotes the effective signal from upstream partner j: on disclosure links, $\tilde{s}_{ij,k,t} = \hat{q}_{j,t}$ with weight $w_{ij,t} = \psi_{ij,t}$; on non-disclosure links, $\tilde{s}_{ij,k,t} = \hat{q}_{j,t} + \hat{\zeta}_{j\ell_j(k),t}$ (the price residual) with weight $w_{ij,t} = \beta^p_{ij,t}$. All weights are partial regression coefficients from the joint projection of $q_{i,t}$ on the full signal vector, and therefore depend on the precision of every signal in the economy.

Effect of partner treatment on signal weights. The SIGE information treatment reduces the private-signal noise $\sigma^2_{\nu,j}$ for treated partner j, making j's posterior $\hat{q}_{j,t}$ more precise. This affects firm i's belief through two channels. First, a more precise $\hat{q}_{j,t}$ raises the residual uncertainty $\gamma_{ij,t} = \mathrm{Var}(\hat{q}_{j,t} \mid \mathcal{I}^{\mathrm{bi},-p}_{k\ell,t})$—the buyer cannot fully anticipate the more informative posterior using other signals—which raises the Kalman weight $\omega_{ij,t}$ in the bilateral pricing problem (equation 19) and the belief weight $\beta^p_{ij,t}$ in the full posterior. Second, the improved precision of j's signals propagates through the input–output structure: partner j's prices and beliefs covary differently with other signals in i's information set, shifting the entire weight vector.

Write the weights under untreated and treated partner precision as $w_{ij}(0)$ and $w_{ij}(1)$, respectively, suppressing time subscripts for clarity. To first order:

$$w_{ij}(T_{j,t}) \;\approx\; w_{ij}(0) \;+\; T_{j,t} \Big[w_{ij}(1) - w_{ij}(0) \Big]. \tag{31}$$

Substituting into (30), the contribution of partner j's signal to firm i's posterior becomes

$$w_{ij}(0) \, \tilde{s}_{ij,k,t} \;+\; T_{j,t} \Big[w_{ij}(1) - w_{ij}(0) \Big] \, \tilde{s}_{ij,k,t}. \tag{32}$$

The first term is a level loading on the partner's signal; the second is the treatment interaction.

Observable signals and the sufficiency of beliefs. The econometrician does not observe the latent signals $\hat{q}_{j,t}$ or $\hat{\zeta}_{j\ell_j(k),t}$ directly. Instead, the regression uses two observable proxies for each partner: the stated expectation $\pi^{e,h}_{j,t}$ (which proxies for $\hat{q}_{j,t}$, j's posterior about the common component) and the price-based signal $g_{j,t}$ (own price growth, which reflects j's marginal cost and therefore loads on both $q_{j,t}$ and $\zeta_{j\ell,t}$).

A key implication of the model is that $\hat{q}_{j,t}$ is a sufficient statistic for the information about $q_{j,t}$ contained in the price residual: on disclosure links, $\beta^p_{ij,t} = 0$ and the buyer ignores the price signal entirely (equation 51). The same logic applies in the regression. Once $\pi^{e,h}_{j,t}$ is included as a regressor, it absorbs the component of partner j's signal that is informative about $q_{j,t}$. The price-based signal $g_{j,t}$, conditional on $\pi^{e,h}_{j,t}$, therefore carries no marginal information about the common component; it captures only the residual variation in partner j's price growth that is orthogonal to j's belief—namely the idiosyncratic component $\hat{\zeta}_{j\ell,t}$ and cost-pass-through effects.

This sufficiency result provides a clean separation of channels in the horse-race specifications (columns 2–5 of Table 3). The $\boldsymbol{\delta}$ coefficients on $\pi^{e,h}_{j,t}$ identify the belief channel—learning about $q_{j,t}$ from partner j's disclosed expectation. The $\boldsymbol{\beta}$ coefficients on $g_{j,t}$, conditional on $\pi^{e,h}_{j,t}$, identify the outcome channel—the direct effect of partner price changes on firm i's expectations through cost pass-through or demand-side adjustments, net of any information about $q_{j,t}$ already captured by beliefs.

Own treatment and the full interaction structure. In the full-sample specification, firm i may also be treated. Reducing $\sigma^2_{\nu,i}$ shifts both κ_i (the weight on i's own private signal, absorbed by the firm fixed effect) and the cross-weights w_{ij} (because i's improved private signal partially substitutes for partner information). A parallel expansion with respect to $T_{i,t}$ adds terms $T_{i,t}\, x_{j,t}$ and $T_{i,t}\, T_{j,t}\, x_{j,t}$, yielding the saturated interaction:

$$f(x_{j,t},\, T_{i,t},\, T_{j,t};\, \boldsymbol{\theta}) \;=\; \theta_0\, x_{j,t} + \theta_1\, T_{j,t}\, x_{j,t} + \theta_2\, T_{i,t}\, x_{j,t} + \theta_3\, T_{i,t}\, T_{j,t}\, x_{j,t}. \tag{33}$$

Mapping coefficients to model objects. Under this derivation, the regression coefficients in equation (6) have the following interpretation:

- $\boldsymbol{\delta}$ coefficients (belief channel, $x_{j,t} = \pi^{e,h}_{j,t}$): δ_0 captures the projection of $\hat{q}_{ik,t}$ on untreated partners' stated beliefs under baseline precision; δ_1 captures the treatment differential—the additional weight on treated partners' beliefs, corresponding to the weight shift $w_{ij}(1) - w_{ij}(0)$ from (31). By sufficiency, this is the channel through which learning about $q_{j,t}$ operates.
- $\boldsymbol{\beta}$ coefficients (outcome channel, $x_{j,t} = g_{j,t}$): conditional on $\pi^{e,h}_{j,t}$, these capture the residual effect of partner price changes—cost pass-through and demand-side adjustments that shift firm i's expectations through channels other than learning about $q_{j,t}$.
- θ_2 coefficients capture the shift in how a treated firm i weights partner signals (own-precision effect); θ_3 captures the interaction between own and partner precision, arising from the curvature of the weights in signal precisions.

Role of fixed effects. The direction×treatment-pair×quarter fixed effects $\alpha_{c(i,j)\times t}$ absorb the level effects of treatment on expectations and, critically, the general-equilibrium price adjustments: when treatment shifts the disclosure set $\mathcal{R}_t$ or the industry-level precision, all firms within a direction–treatment cell are affected symmetrically. The coefficients $\boldsymbol{\beta}$ and $\boldsymbol{\delta}$ are therefore identified from the idiosyncratic variation in partner j's signal—variation due to $\hat{\zeta}_{j\ell,t}$ or to partner-specific shocks—interacted with j's treatment status.

B Expectation disagreement along production links

This appendix documents that the information treatment reduces disagreement in inflation expectations between directly linked buyer–supplier pairs. These results complement the directed learning regressions in Section 3.2 by showing that the treatment improves belief alignment along production links, without assigning a learning role to one side of the relationship.

Measure. For each buyer–supplier link (i, j) active in quarter t, we measure disagreement in h-quarters-ahead inflation expectations as

$$y_{ij,t}^{\text{sym}} \equiv \frac{2\left|\pi_{i,t}^{e,h} - \pi_{j,t}^{e,h}\right|}{\left|\pi_{i,t}^{e,h}\right| + \left|\pi_{j,t}^{e,h}\right|}. \tag{34}$$

This symmetric, scale-free measure captures relative disagreement without mechanically overweighting links involving firms with large absolute forecasts.

Specification. We estimate

$$y_{ij,t}^{\text{sym}} = \alpha_t + \beta_{TU}\,\mathbf{1}\{TU_{ij,t}\} + \beta_{UT}\,\mathbf{1}\{UT_{ij,t}\} + \beta_{TT}\,\mathbf{1}\{TT_{ij,t}\} + X_{ij,t}'\gamma + \varepsilon_{ij,t}, \tag{35}$$

where TU denotes links with a treated buyer and untreated supplier, UT links with an untreated buyer and treated supplier, and TT links in which both firms are treated. The omitted category is untreated–untreated links. Controls $X_{ij,t}$ include sector and geographic fixed effects for both sides of the link and a full interaction of buyer and supplier size categories. All specifications use SIGE sampling weights allocated across links within firm–quarter cells, and cluster standard errors at the firm level.

Average effects. Table B.1 reports results for $h = 6$. Disagreement is significantly lower whenever at least one firm in the link is treated, with the largest reduction on TT links (-0.221, $p < 0.01$). Table B.2 confirms the same pattern at $h = 12$.

Table B.1: Treatment and expectation disagreement ($h = 6$)

	TU	UT	TT
Treatment effect	-0.098^{***}	-0.057^{***}	-0.221^{***}
	(0.020)	(0.016)	(0.019)
Observations		178,996	
R^2		0.241	

Notes: Dependent variable: $y_{ij,t}^{\text{sym}}$ at horizon $h = 6$. TU: treated buyer, untreated supplier. UT: untreated buyer, treated supplier. TT: both treated. Omitted: both untreated. All specifications include quarter fixed effects and baseline dyad controls. Standard errors clustered at the firm level. $^{***}\,p < 0.01$.

Table B.2: Treatment and expectation disagreement ($h = 12$)

	TU	UT	TT
Treatment effect	-0.061^{***}	-0.046^{***}	-0.164^{***}
	(0.020)	(0.016)	(0.019)
Observations		178,996	

Notes: Same as Table B.1 except $h = 12$. $^{***}\,p < 0.01$.

C Proofs

Proof of Proposition 1. **Part (i).** Expand $h^D_{ij,t} = (\mu^D_{ij,t} - 1)(\mu^D_{ij,t})^{-\varepsilon_i(1-\omega_{ij,t})}$ around $\omega_{ij,t} = 0$. At $\omega = 0$: $\mu^D = \mu^*_i$ and $h^D = h(\mu^*_i)$. Differentiating with respect to ω at $\omega = 0$:

$$\left.\frac{\partial h^D}{\partial \omega}\right|_{\omega=0} = \left.\frac{\partial}{\partial \omega}\Big[(\mu^D - 1)(\mu^D)^{-\varepsilon(1-\omega)}\Big]\right|_{\omega=0}.$$

Using $\frac{\partial \mu^D}{\partial \omega}|_{\omega=0} = \varepsilon/(\varepsilon-1)^2 = \mu^*/(varepsilon - 1)$ and $h'(\mu^*) = 0$, the only surviving term is the partial derivative through the exponent:

$$\left.\frac{\partial h^D}{\partial \omega}\right|_{\omega=0} = (\mu^* - 1)(\mu^*)^{-\varepsilon} \cdot \varepsilon \log \mu^* = h(\mu^*_i) \cdot \varepsilon_i \log \mu^*_i \;>\; 0.$$

Hence $h^D_{ij,t} - h(\mu^*_i) = h(\mu^*_i) \cdot \varepsilon_i \log \mu^*_i \cdot \omega_{ij,t} + O(\omega^2)$, which is first order with a strictly positive coefficient.

Part (ii). The seller sets its price on the link to maximize expected profit conditional on its information set. In the Gaussian-linear environment, the belief-dependent component of the optimal price given posterior $\hat{q}_j$ is $\hat{q}_j$ itself (the known components $\bar{p}_{ij}$ and $\hat{\zeta}_{j\ell,t}$ are absorbed into the deterministic part and do not vary with the belief about $q_{j,t}$). The per-period expected profit is

$$\mathbf{E}\Big[\Pi^S_{k\ell,t}\Big(p^*_{k\ell}(\hat{q}_j)\Big)\Big] = h(\mu^*_i)\,\Gamma_{k\ell,t} - \tfrac{1}{2}|h''(\mu^*_i)|\;\mathrm{Var}(q_{j,t} \mid \mathcal{I}_{j\ell})\,\Gamma_{k\ell,t} + O\Big(\mathrm{Var}(q_{j,t} \mid \mathcal{I}_{j\ell})^2\Big),$$

where the second term captures the expected loss from pricing error (the log price deviates from the frictionless target by $(q_{j,t} - \hat{q}_j)$, whose variance is $\mathrm{Var}(q_{j,t} \mid \mathcal{I}_{j\ell})$). The approximation follows from a second-order expansion of the profit function around the frictionless optimum.

With the additional signal, the posterior variance falls by $\Delta\sigma^2_{\text{post}} > 0$. The gain in the seller's expected profit is therefore

$$\Delta\mathbf{E}[\Pi^S] = \tfrac{1}{2}|h''(\mu^*_i)|\,\Delta\sigma^2_{\text{post}}\,\Gamma_{k\ell,t} + O\Big((\Delta\sigma^2_{\text{post}})^2\Big).$$

This is second order in the following precise sense. The reduction in posterior variance from a signal with precision κ_{add} is

$$\Delta\sigma^2_{\text{post}} = \frac{1}{\iota_j} - \frac{1}{\iota_j + \kappa_{\text{add}}} = \frac{\kappa_{\text{add}}}{\iota_j(\iota_j + \kappa_{\text{add}})}, \tag{36}$$

where $\iota_j = 1/\,\mathrm{Var}(q_{j,t} \mid \mathcal{I}_{j\ell})$ is the seller's baseline precision about $q_{j,t}$. The resulting profit gain is

$$\Delta\mathbf{E}[\Pi^S] = \frac{|h''(\mu^*_i)|\,\kappa_{\text{add}}}{2\,\iota_j(\iota_j + \kappa_{\text{add}})}\,\Gamma_{k\ell,t},$$

which is bounded above by $|h''(\mu^*_i)|\Gamma_{k\ell,t}/(2\iota_j^2)$ and vanishes as $\iota_j \to \infty$. This is an *envelope-theorem* effect: because the seller is already optimizing its price given $\mathcal{I}_{j\ell}$, the marginal value of additional information operates only through the second-order reduction in pricing error, not through any first-order change in the markup level.

For the buyer's surplus, a similar expansion gives

$$\mathbf{E}[\Pi^B_{k\ell,t}(p^*_{k\ell})] = \frac{(\mu^*_i)^{1-\varepsilon}}{\varepsilon - 1}\Gamma_{k\ell,t} - \frac{\varepsilon(\mu^*_i)^{-\varepsilon-1}}{2}\;\mathrm{Var}(q_{j,t} \mid \mathcal{I}_{j\ell})\,\Gamma_{k\ell,t} + O(\mathrm{Var}^2).$$

The buyer benefits from the seller's reduced pricing error at rate $\varepsilon(\mu_i^*)^{-\varepsilon-1}/2 \cdot \Delta\sigma^2_{\text{post}}$, which is likewise $O(\Delta\sigma^2_{\text{post}})$—second order relative to the $O(\omega)$ gain from upstream disclosure.

Part (iii). We show that babbling is the unique equilibrium of any cheap-talk game in either direction by establishing a uniformly negative sender bias in each case.

Downstream (buyer → seller). Consider a cheap-talk game in which the buyer sends a message $m \in \mathcal{M}$ about its belief $\hat{q}_{ik,t} \equiv \mathbf{E}_{ik}[q_{i,t}]$ before the seller sets its price.

The buyer's per-period surplus $\Pi^B_{k\ell,t}(\mu) = \mu^{1-\varepsilon_i}\Gamma_{k\ell,t}/(\varepsilon_i - 1)$ is strictly decreasing in the seller's markup μ (and hence in the posted price $p = \mu \cdot mc_{j\ell,t}$) for $\varepsilon_i > 1$. Under the CES specification, the seller's optimal markup $\mu_i^* = \varepsilon_i/(\varepsilon_i - 1)$ is independent of the demand level. However, the seller's posted price $\log p^*_{k\ell,t} = \log \mu_i^* + \hat{q}_{j\ell,t} + \hat{\zeta}_{j\ell,t} + \bar{p}_{ij}$ depends on the seller's posterior $\hat{q}_{j\ell,t}$ about $q_{j,t}$. If the buyer's message m is informative about $q_{j,t}$—as it generally is, since $\text{Cov}(q_{i,t}, q_{j,t}) \neq 0$ through the input-output structure—the seller would update $\hat{q}_{j\ell,t}$ and adjust its price accordingly. For any realization of the buyer's type $\hat{q}_{ik,t}$, the buyer's preferred seller action is the lowest feasible price ($\mu = 1$, yielding $p = \widehat{mc}_{j\ell,t}$), while the seller optimally charges $p = \mu_i^* \cdot \widehat{mc}_{j\ell,t}$ with $\mu_i^* > 1$. The gap $(\mu_i^* - 1) \cdot \widehat{mc}_{j\ell,t} > 0$ is strictly positive for all types and all information sets. By Theorem 1 of Crawford and Sobel (1982), babbling is the unique equilibrium.

Upstream (seller → buyer). Consider a cheap-talk game in which the seller sends a message about its belief $\hat{q}_{j\ell,t} \equiv \mathbf{E}_{j\ell}[q_{j,t}]$ before the buyer sets demand. The seller's profit $\Pi^S_{k\ell,t}(\mu) = (\mu - 1)\mu^{-\varepsilon_i}\Gamma_{k\ell,t}$ is maximized at $\mu_i^* > 1$. By overstating $\hat{q}_{j\ell,t}$, the seller leads the buyer to believe that a higher price is justified by higher costs rather than by a higher markup. Since the buyer's demand depends on the price relative to its perceived outside option $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ (equation (13)), an upward revision in the buyer's estimate of $q_{j,t}$ raises $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$, partially insulating the seller's demand from price increases. The seller therefore strictly prefers to overstate, while the buyer prefers to hear a low cost report (which keeps p low). The bias is again uniformly signed, and babbling is the unique equilibrium by Theorem 1 of Crawford and Sobel (1982).

Variable elasticity. Under demand systems where the elasticity varies with quantity, the seller's optimal markup $\mu(A)$ may depend on the demand shifter $A_{kj,t}$. When $\mu'(A) \geq 0$—the standard case when the demand elasticity is non-increasing in quantity—a higher A raises the markup as well as the price level, widening the gap between the seller's optimal price and the buyer's preferred price and strictly reinforcing the uniform bias in both directions. □

Proof of Proposition 2. Define $R_{ij} \equiv \beta\rho_{ij}/(1 - \beta\rho_{ij})$ and write the establishment condition (26) as a function of $\sigma_q^2 \equiv \sigma^2_{q,j}$, holding all other link characteristics fixed.

Step 1: Scaling of the cooperative surplus. Write $W'_{k\ell,t}(\mu) = w'(\mu)\,\Gamma_{k\ell,t}$ where $w'(\mu) \equiv \varepsilon\mu^{-\varepsilon-1}(1 - \mu)$, so that from (23),

$$\Delta W_{k\ell,t} = |w'(\mu_i^*)| \cdot \frac{\varepsilon\,\omega_{ij,t}}{(\varepsilon - 1)^2}\,\Gamma_{k\ell,t} + O(\omega^2).$$

For small σ_q^2, the informational content of the price satisfies $\omega_{ij,t} \approx \gamma_{ij,t}/\sigma^2_{\zeta,j}$ (from (19) for small $\gamma_{ij,t}$). Since $\gamma_{ij,t} \leq \text{Var}(\hat{q}_{j\ell,t}) = \sigma^4_{q,j}/(\sigma^2_{q,j} + \sigma^2_{\nu,j}) \leq \sigma^2_{q,j}$, we have $\omega_{ij,t} = O(\sigma_q^2/\sigma^2_{\zeta,j})$. Hence

$$\Delta W_{k\ell,t} = \underbrace{\frac{\varepsilon\,|w'(\mu_i^*)|}{(\varepsilon - 1)^2\sigma^2_{\zeta,j}}}_{\equiv\ c_W > 0}\ \sigma_q^2\,\Gamma_{k\ell,t}\ +\ O\big(\sigma_q^4\big)\,. \tag{37}$$

Step 2: Threshold from the disclosure IC. The IC (26) requires

$$f \;\leq\; \Pi^{\text{net}}_{ij,t}\,\Gamma_{k\ell,t},$$

where

$$\Pi^{\text{net}}_{ij,t} \;\equiv\; [h(\mu_i^*) - h^D_{ij,t}] + R_{ij}[h(\mu_i^*) - h(\mu^D_{ij,t})].$$

At $\omega_{ij,t} = 0$: $\mu^D_{ij,t} = \mu_i^*$, so $h^D_{ij,t} = h(\mu_i^*) = h(\mu^D_{ij,t})$ and $\Pi^{\text{net}}_{ij,t} = 0$. For $\omega_{ij,t} > 0$, $h^D_{ij,t} > h(\mu_i^*) > h(\mu^D_{ij,t})$, and using $\omega_{ij,t} = O(\sigma_q^2/\sigma^2_{\zeta,j})$:

$$\Pi^{\text{net}}_{ij,t} \;=\; \underbrace{\frac{\varepsilon_i\,|w'(\mu_i^*)|}{(\varepsilon_i - 1)^2\sigma^2_{\zeta,j}}}_{\equiv\; c_\Pi\,>\,0}(R_{ij} - 1)\,\sigma_q^2 \;+\; O(\sigma_q^4). \tag{38}$$

Hence $\Pi^{\text{net}}_{ij,t} > 0$ whenever $R_{ij} > 1$, i.e. $\beta\rho_{ij} > 1/2$, which holds for all links in our calibration. Substituting and dividing by $\Gamma_{k\ell,t} > 0$:

$$\frac{f}{\Gamma_{k\ell,t}} \;\leq\; (R_{ij} - 1)\,c_\Pi\,\sigma_q^2 \;+\; O(\sigma_q^4).$$

The left-hand side is a positive constant; the right-hand side equals zero at $\sigma_q^2 = 0$ and increases without bound as σ_q^2 grows. By the intermediate value theorem there exists a unique threshold $\bar\sigma^2_{q,ij} > 0$. Disclosure is established if and only if $\sigma_q^2 \geq \bar\sigma^2_{q,ij}$.

To leading order, the threshold solves

$$\frac{f}{\Gamma_{k\ell,t}} \;=\; (R_{ij} - 1)\,c_\Pi\,\bar\sigma^2_{q,ij}, \tag{39}$$

giving the closed-form expression $\bar\sigma^2_{q,ij} = f/[(R_{ij} - 1)\,c_\Pi\,\Gamma_{k\ell,t}]$.

Step 3: Comparative statics. From the closed-form threshold, each comparative static follows:

(a) *Decreasing in ρ_{ij}.* The numerator f is constant; the denominator $(R_{ij}-1)\,c_\Pi\,\Gamma$ is increasing in R_{ij} and hence in ρ_{ij}. Hence $\bar\sigma^2_{q,ij}$ is decreasing in ρ_{ij}.

(b) *Decreasing in $\Gamma_{k\ell,t}$.* The denominator contains $\Gamma_{k\ell,t}$ directly, so the threshold is decreasing in trade volume.

(c) *Increasing in supplier breadth.* Higher breadth raises $\sigma^2_{\zeta,j}$, which reduces c_Π and hence raises the threshold.

The discount factor β enters R_{ij} multiplicatively: higher β raises R_{ij} and uniformly lowers $\bar\sigma^2_{q,ij}$, but since β is common across links it does not generate cross-link heterogeneity. □

D Microfoundation of bilateral demand

This appendix microfounds the CES demand specification (13) used in Section 5.1 from a model in which the buyer can switch to a competitive seller in the supplier's industry, generating both an intensive and an extensive margin of demand.

Notation. Consider a bilateral link between supplier ℓ in industry j and buyer k in industry i, as in Section 5.1. Let $\varepsilon_i > 1$ denote the bilateral demand elasticity, $Q_{j,t}$ the average cost of production in the supplier's industry j (the price the buyer would pay if it switched to a competitive seller in j), $S_{i,t}$ the buyer's total nominal expenditure on inputs, and a_{ij} the input-output weight (share of industry i's expenditure allocated to industry j's input). The seller sets a bilateral price p. The buyer's information set is $\mathcal{I}_{k,t}$.

D.1 Expenditure shares and intensive margin

Conditional on sourcing from supplier ℓ, buyer k allocates a constant expenditure share a_{ij} to industry j's input. On the intensive margin, the quantity demanded at bilateral price p is

$$x^{\text{int}}(p) \;=\; \frac{a_{ij}\, S_{i,t}}{p}, \tag{40}$$

yielding a unit price elasticity. This is the standard result under Cobb–Douglas aggregation over inputs at the industry level.

D.2 Switching cost and extensive margin

The buyer has access to an outside option: it can purchase industry j's input from a competitive seller at the industry-average cost of production $Q_{j,t}$. Switching entails a match-specific cost $\tau\, Q_{j,t}$, where $\tau \geq 1$ is a random switching cost drawn with cumulative distribution $G(\cdot)$ and density $g(\cdot)$. The buyer observes the realization of τ before making its sourcing decision; the seller does not, and sets p without knowledge of τ. Crucially, the buyer does not observe $Q_{j,t}$ directly. As established in the main text, $\log Q_{j,t} = \bar{Q}_j + q_{j,t}$, so uncertainty about $Q_{j,t}$ is equivalent to uncertainty about the common component $q_{j,t}$. The buyer evaluates the outside option at the conditional expectation $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$.

The buyer sources from supplier ℓ if

$$p \;\leq\; \tau\, \mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]\,. \tag{41}$$

From the seller's perspective, the probability that the buyer does not switch is the extensive margin:

$$\Pr(\text{source from } \ell) = 1 - G\!\left(\frac{p}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}\right). \tag{42}$$

D.3 Aggregate bilateral demand

Combining the intensive margin (40) with the extensive margin (42), the bilateral demand at price p is

$$x(p) = \frac{a_{ij}\, S_{i,t}}{p}\left[1 - G\!\left(\frac{p}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}\right)\right]. \tag{43}$$

D.4 CES specification

The bilateral demand (43) has constant price elasticity ε_i if and only if the survival function $1 - G$ is a power law:

$$1 - G(\tau) = \tau^{-(\varepsilon_i - 1)}, \qquad \tau \geq 1,\ \varepsilon_i > 1. \tag{44}$$

This is a Pareto distribution with shape parameter $\varepsilon_i - 1$ and support $[1, \infty)$. The restriction $\varepsilon_i > 1$ ensures finite mean switching costs; higher ε_i concentrates the distribution near $\tau = 1$, meaning the buyer is nearly indifferent and demand is highly elastic.

Under (44), equation (43) reduces to the CES form.

Derivation. Let $r \equiv p/\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$. Substituting (44) into (43):

$$x(p) = \frac{a_{ij}\, S_{i,t}}{p}\; r^{-(\varepsilon_i - 1)} = \frac{a_{ij}\, S_{i,t}}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}\; \left(\frac{p}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}\right)^{-\varepsilon_i}.$$

□

This yields the CES demand

$$x_{kj,t}(p) \;=\; \left(\frac{p}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}\right)^{-\varepsilon_i} A_{kj,t}, \qquad A_{kj,t} \;\equiv\; \frac{a_{ij}\, S_{i,t}}{\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}, \tag{45}$$

which is equation (13). The total price elasticity is ε_i: the extensive margin (switching) contributes $\varepsilon_i - 1$, since $\partial \log[1 - G(r)]/\partial \log r = -(\varepsilon_i - 1)$ from (44); the intensive margin contributes one unit from (40).

Connection to the markup distortion. The conditional expectation $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ embeds the buyer's inference problem about the supplier industry's average cost of production. Since $\log Q_{j,t} = \bar{Q}_j + q_{j,t}$, learning about $Q_{j,t}$ is equivalent to learning about $q_{j,t}$, the common component of industry j's marginal costs.

On *disclosure links*, the buyer observes the supplier's belief $\hat{q}_{j,t}$ directly, so $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ does not depend on p and the demand elasticity is constant at ε_i; the seller charges the frictionless markup μ_i^*.

On *non-disclosure links*, the buyer uses p to update beliefs about $q_{j,t}$, so the effective demand response to p incorporates the signal-extraction weight $\omega_{ji,t}$. The perceived elasticity becomes $\varepsilon_i(1 - \omega_{ji,t})$, and the markup distortion (21) follows.

E Surplus derivations and breadth–noise condition

Bilateral surplus derivations

Using the CES demand (13), the supplier's per-period profit on the link between buyer k and supplier ℓ is

$$\Pi^S_{k\ell,t}(\mu) = (p - mc_{j\ell,t})\, x_{kj,t} = (\mu - 1)\mu^{-\varepsilon}\Gamma_{k\ell,t} \equiv h(\mu)\, \Gamma_{k\ell,t}. \tag{46}$$

The buyer's surplus is the area under the demand curve above the transaction price:

$$\Pi^B_{k\ell,t} = \int_p^\infty x_{kj,t}(t)\, dt = \frac{A_{kj,t}}{\varepsilon - 1}\, p^{1-\varepsilon},$$

where the integral converges because $\varepsilon > 1$. Substituting $p = \mu\, mc_{j\ell,t}$ and using $\Gamma_{k\ell,t} \equiv mc_{j\ell,t}^{1-\varepsilon} A_{kj,t}$ gives

$$\Pi^B_{k\ell,t}(\mu) = \frac{1}{\varepsilon - 1}\, \mu^{1-\varepsilon}\, \Gamma_{k\ell,t}. \tag{47}$$

Total bilateral surplus is

$$W_{k\ell,t}(\mu) \equiv \Pi^S_{k\ell,t}(\mu) + \Pi^B_{k\ell,t}(\mu) = \left[(\mu-1)\mu^{-\varepsilon} + \frac{\mu^{1-\varepsilon}}{\varepsilon-1}\right]\Gamma_{k\ell,t}. \tag{48}$$

Differentiating:

$$W'_{k\ell,t}(\mu) = \varepsilon\,\mu^{-\varepsilon-1}(1-\mu)\,\Gamma_{k\ell,t}. \tag{49}$$

Total surplus is maximized at $\mu = 1$ (marginal cost pricing) and is strictly decreasing for $\mu > 1$. The envelope conditions at μ_i^* are:

$$h'(\mu_i^*) = 0, \qquad h''(\mu_i^*) < 0. \tag{50}$$

Breadth–noise scaling

Supplier breadth has *competing* effects on $\omega_{ij,t}$ (19). Higher supplier breadth increases $\sigma^2_{\zeta,j}$ (denominator), which pushes ω down. But it simultaneously increases the buyer's *residual uncertainty* $\gamma_{ij,t}$ (numerator): a broad supplier has $q_{j,t} = \sum_r h_{jr}\alpha_r(\log w_t - \bar{w}) + \sum_r h_{jr}e_{r,t}$ from (10), which loads on many diverse industry-specific shocks that the buyer cannot easily learn from its private signal $s_{q,k,t}$ or from the upstream attention signal. The net effect is that $\omega_{ij,t}$ increases with supplier breadth under the following condition.

Assumption 2 (Breadth–noise scaling)**.** *Let $B_j \equiv |\{r : h_{jr} > 0,\ r \neq j\}|$ denote the effective upstream breadth of industry j (the number of industries whose shocks enter $q_{j,t}$ through the Leontief inverse). The industry-specific shocks $\{e_{r,t}\}_{r=1}^n$ have uniformly bounded variances $\underline{\sigma}_e^2 \leq \mathrm{Var}(e_{r,t}) \leq \bar{\sigma}_e^2$. Then for all j and all buyers k with supplier ℓ in industry j:*

$$\gamma_{ij,t} \equiv \mathrm{Var}\left(\hat{q}_{j\ell,t} \mid \mathcal{I}^{\mathrm{bi},-p}_{k\ell,t}\right) \;\geq\; c\,B_j\,\underline{\sigma}_e^2 \qquad \textit{for some } c > 0.$$

Assumption 2 holds when the Leontief weights h_{jr} spread across many industries with independent shocks $e_{r,t}$: the common component $q_{j,t} = \sum_r h_{jr}e_{r,t} + \cdots$ becomes a weighted average of many independent terms, each of which is hard for the buyer to learn individually from its private signal about $q_{i,t}$ (which loads on a different set of Leontief weights h_{ir}). Meanwhile, $\sigma^2_{\zeta,j}$ grows at most proportionally to B_j (since the idiosyncratic component (12) aggregates B_j independent firm-specific shocks, each with bounded variance). Under Assumption 2, the numerator of (19) grows faster than the denominator in B_j, so $\omega_{ij,t}$ is increasing in supplier breadth.

The implication is that, on non-disclosure links, *price-based learning becomes more important for broad suppliers*: the buyer has no better window into the supplier's costs than the price itself.[37]

F Network equilibrium derivations

This appendix provides full derivations supporting Section 5. We collect (i) the complete signal-extraction system at the firm level, (ii) the aggregation to industry-level average beliefs, (iii) the matrix characterization of the equilibrium, and (iv) the definition and fixed-point properties of the disclosure set $\mathcal{R}_t$.

[37] When non-cooperative signals are introduced (Appendix F and Appendix G), the assumption acquires an additional term $-\iota^{-1}_{\mathrm{up},k\ell}$ on the right-hand side, reflecting the precision of the non-cooperative signal. The qualitative conclusion is unchanged: the bound requires only that the residual uncertainty grows faster than the noise.

F.1 Signal extraction at the firm level

Consider firm k in industry i. The idiosyncratic prior $\hat{\zeta}_{ik,t}$ is formed in stage (i) of the timing protocol (Section 5.1) from the signals $z_{k,t}$ and $\{s_{z,k\ell,t}\}$ and is held fixed thereafter. The signal extraction problem below therefore conditions on $\hat{\zeta}_{ik,t}$ and concerns only the common component $q_{i,t}$.

In period t, firm k's information set for the common component consists of:

(a) the private signal $s_{q,k,t} = q_{i,t} + \nu_{k,t}$ with precision $\iota_{\nu,i} \equiv 1/\sigma^2_{\nu,i}$;

(b) for each upstream partner $j \in \mathcal{N}^{\text{up}}_{ik}$, a non-cooperative signal $s^{\text{up}}_{ji,k,t} = \hat{q}_{j,t} + \eta^{\text{up}}_{ji,t}$ with precision $\iota_{\text{up},ji} \equiv 1/\sigma^2_{\text{up},ji}$;

(c) for each downstream partner $\ell \in \mathcal{N}^{\text{dn}}_{ik}$, a non-cooperative signal $s^{\text{dn}}_{i\ell,k,t} = \hat{q}_{\ell,t} + \eta^{\text{dn}}_{i\ell,t}$ with precision $\iota_{\text{dn},i\ell} \equiv 1/\sigma^2_{\text{dn},i\ell}$;

(d) for each disclosure link $(j,i) \in \mathcal{R}_t$, the supplier's posterior $\hat{q}_{j,t}$ (observed without noise);

(e) for each non-disclosure link $(j,i) \notin \mathcal{R}_t$, the price residual $\hat{q}_{j,t} + \hat{\zeta}_{j\ell_j(k),t}$ with Kalman weight $\omega_{ij,t}$ from equation (19).

In the Gaussian-linear environment, the posterior is the multivariate regression of $q_{i,t}$ on the full signal vector available to the buyer:

$$\hat{q}_{ik,t} = \kappa_i\, s_{q,k,t} \;+\; \sum_{j:\,(j,i)\in\mathcal{R}_t} \psi_{ij,t}\, \hat{q}_{j,t} \;+\; \sum_{j:\,(j,i)\notin\mathcal{R}_t} \beta^{p}_{ij,t} \left(\hat{q}_{j,t} + \hat{\zeta}_{j\ell_j(k),t}\right) + \sum_{j\in\mathcal{N}^{\text{up}}_{ik}} \beta^{\text{up}}_{ij,t}\, s^{\text{up}}_{ji,k,t} \;+\; \sum_{\ell\in\mathcal{N}^{\text{dn}}_{ik}} \phi_{i\ell,t}\, s^{\text{dn}}_{i\ell,k,t}, \tag{51}$$

where the weights $(\kappa_i, \psi_{ij,t}, \beta^{p}_{ij,t}, \beta^{\text{up}}_{ij,t}, \phi_{i\ell,t})$ are the partial regression coefficients from the joint regression of $q_{i,t}$ on all signals available to buyer k in industry i. The multivariate regression accounts for the cross-correlation between different suppliers' price residuals (through the Leontief inverse H): a supplier whose $q_{j,t}$ is highly correlated with $q_{j',t}$ of another supplier receives a lower partial coefficient because it is partially redundant. The weight $\beta^{p}_{ij,t}$ on the price residual from a non-disclosure supplier j is the belief system's coefficient—distinct from the bilateral $\omega_{ij,t}$ of (19), which governs the markup distortion.[38] On disclosure links, sufficiency of the disclosed belief implies $\beta^{p}_{ij,t} = 0$ and $\beta^{\text{up}}_{ij,t} = 0$: the buyer ignores both the price residual and the non-cooperative upstream signal for that supplier.

F.2 Aggregation to industry-level beliefs

Integrating (51) over $k \in [0,1]$ and applying the law of large numbers (the idiosyncratic noise $\hat{\zeta}_{j\ell_j(k),t}$, $\eta^{\text{up}}_{ji,t}$, and $\nu_{k,t}$ are independent across the different partners that different firms k face), the industry-average belief satisfies equation (52) in the main text:

$$\bar{\hat{q}}_{i,t} = \kappa_i\, q_{i,t} + \sum_{j:\,(j,i)\in\mathcal{R}_t} \psi_{ij,t}\, \bar{\hat{q}}_{j,t} + \sum_{j:\,(j,i)\notin\mathcal{R}_t} \beta^{p}_{ij,t}\, q_{j,t} + \sum_{\ell\in\mathcal{N}^{\text{dn}}_{i}} \phi_{i\ell,t}\, \bar{\hat{q}}_{\ell,t} + \eta_{i,t}, \tag{52}$$

where $\eta_{i,t} = O(1/\sqrt{N_i})$ collects residual averaging noise (vanishing in the continuum limit).

[38] The bilateral $\omega_{ij,t}$ is the Kalman weight the *seller* internalizes when setting the price: it reflects the buyer's sensitivity at the bilateral stage, controlling only for the private signal and the upstream attention signal. The belief weight $\beta^{p}_{ij,t}$ reflects the price signal's contribution to the full posterior, after controlling for all other signals.

F.3 Matrix characterization

Stacking across n industries, define:

- Ψ_t^R: the $n \times n$ matrix with (i,j)-entry $\psi_{ij,t}\,\mathbf{1}[(j,i) \in \mathcal{R}_t]$ (disclosed-belief weights, upstream);
- Φ_t: the $n \times n$ matrix with (i,ℓ)-entry $\phi_{i\ell,t}$ (downstream non-cooperative weights);
- $\kappa \equiv \mathrm{diag}(\kappa_1, \ldots, \kappa_n)$ (private-signal weights);
- B_t^{NR}: the $n \times n$ matrix with (i,j)-entry $\beta_{ij,t}^{p}\,\alpha_j\,\mathbf{1}[(j,i) \notin \mathcal{R}_t]$ (price-channel contribution to exogenous loading on $q_{j,t}$).

Define $B_t^{\text{exog}} \equiv \kappa + B_t^{NR}$ (the exogenous block: weights on fundamental $\mathbf{q}_t$) and $B_t^{\text{endo}} \equiv \Psi_t^R + \Omega_t^{\text{up}} + \Phi_t$ (the endogenous block: weights on other industries' beliefs).[39] The system (53) is

$$\bar{\hat{\mathbf{q}}}_t = B_t^{\text{endo}}\,\bar{\hat{\mathbf{q}}}_t + B_t^{\text{exog}}\,\mathbf{q}_t + \boldsymbol{\eta}_t. \tag{53}$$

Provided the spectral radius $\rho(B_t^{\text{endo}}) < 1$—which holds because the endogenous weights are partial regression coefficients that absorb only a fraction of the total signal precision—the unique solution is (54):

$$\bar{\hat{\mathbf{q}}}_t = [I - B_t^{\text{endo}}]^{-1}\,B_t^{\text{exog}}\,\mathbf{q}_t + [I - B_t^{\text{endo}}]^{-1}\boldsymbol{\eta}_t. \tag{54}$$

The matrix $\mathcal{H}_t \equiv [I - B_t^{\text{endo}}]^{-1}$ is the *informational Leontief inverse*. Its (i,j)-entry captures the total (direct plus indirect) effect of a marginal change in the information flowing into industry j on industry i's average belief, through all paths in the information-transmission network. The belief matrix $M_t \equiv \mathcal{H}_t\,B_t^{\text{exog}}$ maps fundamentals $\mathbf{q}_t$ into beliefs $\bar{\hat{\mathbf{q}}}_t$; its diagonal entry $M_{ii,t}$ is the accuracy of industry i's average belief about its own common component.

F.4 Disclosure set and fixed point

The disclosure set $\mathcal{R}_t$ is defined as the largest set of links satisfying the establishment condition (26) simultaneously:

$$\mathcal{R}_t \equiv \Big\{(j,i):\ a_{ij} > 0 \ \text{ and } \ f \le R_{ij}\,\Delta W_{ji,t}\Big\},$$

where $R_{ij} \equiv \beta\rho_{ij}/(1-\beta\rho_{ij})$ capitalizes per-period flows over the expected match life, f is the uniform setup cost (see Section 5.3), and $\Delta W_{ji,t}$ is the cooperative surplus (23). The cooperative surplus depends on the bilateral markup weight $\omega_{ij,t}$ from (19), which in turn depends on the disclosure configuration $\mathcal{R}_t$ (through the attention signal precision). The equilibrium is therefore a fixed point.

[39] Ω_t^{up} collects the upstream attention weights $\beta_{ij,t}^{\text{up}}\,\mathbf{1}[(j,i) \notin \mathcal{R}_t]$. In the aggregated equation it multiplies $\bar{\hat{q}}_{j,t}$, since the attention signal's mean component is the seller's average belief.

Computation. We solve by iterating on $\mathcal{R}_t$. Starting from the empty set $\mathcal{R}_t^{(0)} = \varnothing$ (all links non-disclosure), each iteration:

1. Computes the bilateral weights $\omega_{ij,t}^{(s)}$ and the belief system weights $(B_t^{\text{exog},(s)}, B_t^{\text{endo},(s)})$ given $\mathcal{R}_t^{(s)}$.
2. Evaluates $\Delta W_{ji,t}^{(s)}$ for each active link.
3. Updates $\mathcal{R}_t^{(s+1)} = \{(j,i) : f \leq R_{ij}\, \Delta W_{ji,t}^{(s)}\}$.

Because adding a link to $\mathcal{R}_t$ reduces the markup distortion on that link (lowering $\omega_{ij,t}$) and weakly relaxes the establishment condition on neighboring links (through the informational Leontief inverse), the mapping $\mathcal{R}_t^{(s)} \mapsto \mathcal{R}_t^{(s+1)}$ is monotone on the lattice of subsets of feasible links. By Tarski's fixed-point theorem, a greatest fixed point exists and is reached by the iteration from above (starting from the full set and pruning links that violate the establishment condition). Iteration from below (starting from the empty set) converges to the least fixed point. In all numerical exercises in this paper, both iterations converge to the same set, indicating uniqueness.

G Non-cooperative signals and endogenous attention

This appendix extends the baseline model of Section 5.1 by introducing non-cooperative signals—noisy observations of trading partners' beliefs that are generated passively on every link—and endogenizing their precision through a costly attention problem. In the baseline model, the buyer's only sources of information about the supplier's common component $q_{j,t}$ are the private signal $s_{q,k,t}$ and the bilateral price residual (18). The non-cooperative signal adds a third source that does not depend on the disclosure decision, generating a channel of information transmission that is insensitive to relationship characteristics—matching the downstream empirical patterns documented in Section 4.

G.1 Non-cooperative information acquisition

Both the buyer and the supplier can acquire information about each other from observable features of their trading partners' network activity. The distinguishing feature of these signals is that they do not require the observed party's voluntary cooperation: the receiver may actively invest in attention, but the information is extracted from observable activity rather than from a deliberate message. Neither party has an incentive to block the other's observation, since both benefit from more accurate pricing.

Upstream signal. Buyer k (in industry i) can observe noisy indicators of supplier ℓ's (in industry j) activity *beyond* the bilateral price that ℓ charges to k. These include capacity utilization, delivery lead times, order backlogs, and inventory movements—all of which are correlated with the supplier's belief about aggregate conditions but are not generated by the pricing decision.

Let $a_{j\ell,t}$ denote a summary statistic of supplier ℓ's observable non-price activity. In the Gaussian-linear environment, $a_{j\ell,t}$ depends on ℓ's belief and on an independent activity shock:

$$a_{j\ell,t} = \lambda_j \hat{q}_{j\ell,t} + u_{j\ell,t}, \qquad u_{j\ell,t} \sim \mathcal{N}(0, \sigma_{u,j}^2),$$

where $\lambda_j > 0$ measures the sensitivity of industry j's activity to beliefs and $u_{j\ell,t}$ is independent of $(q_{j,t}, \zeta_{j\ell,t})$. Adding observation noise $\nu^{\text{up}}_{k\ell,t} \sim \mathcal{N}(0, \sigma^2_{\nu,k\ell})$ to capture imperfect attention and normalizing by λ_j, the buyer extracts the signal

$$s^{\text{up}}_{k\ell,t} \equiv \frac{a_{j\ell,t} + \nu^{\text{up}}_{k\ell,t}}{\lambda_j} = \hat{q}_{j\ell,t} + \underbrace{\frac{u_{j\ell,t} + \nu^{\text{up}}_{k\ell,t}}{\lambda_j}}_{\equiv\ \eta^{\text{up}}_{k\ell,t}}, \tag{55}$$

with noise variance

$$\sigma^2_{\text{up},k\ell} = \frac{\sigma^2_{u,j} + \sigma^2_{\nu,k\ell}}{\lambda_j^2}. \tag{56}$$

Because $u_{j\ell,t}$ is independent of the seller's idiosyncratic component $\hat{\zeta}_{j\ell,t}$ (which depends on productivity shocks, not activity shocks), the non-cooperative signal $s^{\text{up}}_{k\ell,t}$ and the bilateral price residual (18) are conditionally independent given $\hat{q}_{j\ell,t}$. The two sources therefore provide *separate* information about the supplier's belief.

When the non-cooperative signal is included, buyer k's information set on the link with supplier ℓ expands to $\mathcal{I}^{\text{bi}}_{k\ell,t} = \{s_{q,k,t},\ s^{\text{up}}_{k\ell,t},\ \text{price residual}\}$, and the residual uncertainty $\gamma_{ij,t}$ in (19) is now conditioned on both the private signal *and* the non-cooperative signal, which is lower than in the baseline. The non-cooperative signal therefore *reduces* $\omega_{ij,t}$, shrinking the markup distortion and the cooperative surplus from disclosure.

Downstream signal. Seller ℓ (in industry j) can observe noisy indicators of buyer k's (in industry i) activity—principally buyer k's output prices (charged to k's own downstream customers), but also order patterns with other suppliers and inventory movements. From (16), each of buyer k's output prices has the form

$$\log p^*_{kk',t} = \bar{p}_{i'i} + \hat{q}_{ik,t} + \hat{\zeta}_{ik,t},$$

where i' denotes the industry of downstream buyer k' and $\bar{p}_{i'i}$ is the steady-state bilateral price (which depends only on the industry pair), and $\hat{\zeta}_{ik,t}$ is the *same* across all of k's output links (it depends on k's upstream matches, not on the identity of the downstream buyer k'). After subtracting the known component $\bar{p}_{i'i}$, every output price yields the same residual $\hat{q}_{ik,t} + \hat{\zeta}_{ik,t}$.

Adding observation noise $\nu^{\text{dn}}_{\ell k,t} \sim \mathcal{N}(0, \sigma^2_{\nu,\ell k})$ (reflecting imperfect attention), the seller extracts

$$s^{\text{dn}}_{\ell k,t} = \hat{q}_{ik,t} + \underbrace{\hat{\zeta}_{ik,t} + \nu^{\text{dn}}_{\ell k,t}}_{\equiv\ \eta^{\text{dn}}_{\ell k,t}}, \tag{57}$$

with noise variance

$$\sigma^2_{\text{dn},\ell k} = \sigma^2_{\zeta,i} + \sigma^2_{\nu,\ell k}, \tag{58}$$

where $\sigma^2_{\zeta,i} \equiv \text{Var}(\hat{\zeta}_{ik,t} \mid \mathcal{I}_{j\ell,t})$ is the variance of the buyer's *posterior* idiosyncratic component as seen from the seller's information set. Because $\hat{\zeta}_{ik,t}$ is common across all of k's output links, observing multiple output prices does not reduce the idiosyncratic noise—it is a firm-level confound that cannot be averaged out. The signal $s^{\text{dn}}_{\ell k,t}$ is informative about $q_{j,t}$ because $q_{i,t}$ and $q_{j,t}$ are correlated through the common underlying shocks.

Since the buyer's idiosyncratic component $\hat{\zeta}_{ik,t}$ depends on k's upstream matches and the seller's idiosyncratic component $\hat{\zeta}_{j\ell,t}$ depends on ℓ's upstream matches, and these match sets are disjoint (zero-overlap property), the downstream signal $s^{\text{dn}}_{\ell k,t}$ is conditionally independent of the bilateral price residual (18) given $\hat{q}_{ik,t}$ and $\hat{q}_{j\ell,t}$.

Effect on the baseline model. The non-cooperative signals modify two objects in the baseline. First, the Kalman weight $\omega_{ij,t}$ on the price residual falls (the buyer has an additional source of information about $q_{j,t}$), reducing the markup distortion and the cooperative surplus from disclosure. Second, the seller receives a noisy signal $s^{\text{dn}}_{\ell k,t}$ about the buyer's belief, generating a positive but second-order downstream information channel that does not require voluntary disclosure. This is the channel that maps to the empirical finding of positive downstream transmission that is insensitive to relationship characteristics.

G.2 Attention allocation across links

The observation noise $\sigma^2_{\nu,k\ell}$ in both the upstream and downstream signals is not a technological primitive—it reflects the *attention* that firm k devotes to the link with ℓ. Tracking a supplier's lead times, parsing a buyer's order patterns, or following industry news all require scarce managerial time and organizational bandwidth. We model this as a costly attention problem: each firm chooses how much attention to devote to each trading partner, facing a convex cost of attention, and the precision of non-cooperative signals responds endogenously to this choice.

Attention technology. Firm k in industry i allocates attention $\lambda^{\text{att}}_{k\ell} \geq 0$ to each upstream link $\ell \in \mathcal{N}^{\text{up}}_{ik}$ and $\lambda^{\text{att}}_{\ell k} \geq 0$ to each downstream link $k' \in \mathcal{N}^{\text{dn}}_{ik}$. The observation noise on the link between k and ℓ is decreasing in attention:

$$\sigma^2_{\nu,k\ell} = \frac{\sigma^2_0}{\lambda^{\text{att}}_{k\ell}}, \tag{59}$$

where $\sigma^2_0 > 0$ is a baseline noise parameter common across links. Attention is costly: each unit allocated to the link has convex cost $c(\lambda^{\text{att}}_{k\ell})$, where c is twice differentiable with $c(0) = 0$, $c' > 0$, and $c'' > 0$.[40]

Timing and equilibrium concept. Both attention and disclosure are determined within each period, but at different stages. The within-period timing extends the baseline protocol:

(0a) The economy is in steady state with industry-level uncertainty $\sigma^2_{q,j}$. Counterfactual exercises consider unforeseen (MIT) shocks to $\sigma^2_{q,j}$, after which firms re-optimize attention and disclosure given the new level.[41]

(0b) Each firm k chooses attention allocation $\{\lambda^{\text{att}}_{k\ell}\}$ across links, given the uncertainty regime and anticipating the equilibrium disclosure set.

[40] A natural parametrization is $c(\lambda) = \frac{\gamma_i}{2}\lambda^2$, where γ_i varies across industries with organizational capacity. The quadratic case yields closed-form solutions; the results below hold for any strictly convex cost.

[41] Firms need not observe the *realization* of $q_{j,t}$—only the *variance*, which captures the aggregate uncertainty environment.

(i)–(iv) The period unfolds as described in the timing protocol of Section 5.1, with the addition that non-cooperative signals $s^{\text{up}}_{k\ell,t}$ and $s^{\text{dn}}_{\ell k,t}$ are observed on every link in stage (i).

The equilibrium is a fixed point in attention and disclosure conditional on the uncertainty regime:

$$\Big(\{\lambda^{\text{att},*}_{k\ell,t}\},\ \mathcal{R}_t\Big) \quad \text{such that} \quad \begin{cases} \lambda^{\text{att},*}_{k\ell,t} \text{ is optimal given } \mathcal{R}_t \text{ and } \sigma^2_{q,j}, \\ \mathcal{R}_t \text{ satisfies (26) given } \sigma^2_{\nu,k\ell}(\lambda^{\text{att},*}_{k\ell,t}). \end{cases}$$

Optimal attention. The firm's problem on each upstream link is

$$\max_{\lambda^{\text{att}}_{k\ell}\geq 0}\ a_{ij}\,\Gamma_{k\ell,t}\,g\Big(\omega_{ij,t}(\lambda^{\text{att}}_{k\ell},\sigma^2_{q,j})\Big)\ -\ c\Big(\lambda^{\text{att}}_{k\ell}\Big), \tag{60}$$

where $g(\omega)$ is the per-unit surplus gain from reducing the markup distortion. The first-order condition equates the marginal value of attention to its marginal cost:

$$a_{ij}\,\Gamma_{k\ell,t}\,g'(\omega)\,\frac{\partial\omega}{\partial\lambda^{\text{att}}_{k\ell}}\ =\ c'(\lambda^{\text{att},*}_{k\ell,t}). \tag{61}$$

Under the quadratic cost $c(\lambda)=\frac{\gamma_i}{2}\lambda^2$, the optimal attention is

$$\lambda^{\text{att},*}_{k\ell,t}\ =\ \frac{1}{\gamma_i}\,a_{ij}\,\Gamma_{k\ell,t}\,\mathcal{V}_{k\ell,t}(\sigma^2_{q,j}), \tag{62}$$

where $\mathcal{V}_{k\ell,t}(\sigma^2_{q,j})\equiv g'(\omega)\cdot|\partial\omega/\partial\lambda^{\text{att}}_{k\ell}|$ captures the marginal informational value.

Properties of optimal attention. Four properties of (62) generate testable predictions:

1. **Trade volume.** Firms devote more attention to larger trading partners. A supplier representing a larger share of the buyer's costs (a_{ij} high) or a larger bilateral flow ($\Gamma_{k\ell,t}$ high) receives more precise non-cooperative signals.

2. **Uncertainty.** Attention on each link is increasing in $\sigma^2_{q,j}$: when aggregate uncertainty is high, the markup distortion is more costly and the marginal value of attention rises. Following an MIT shock that raises $\sigma^2_{q,j}$, *both* margins activate simultaneously: disclosure links expand (Proposition 2) *and* non-cooperative signals on non-disclosure links become more informative.

3. **Substitution with disclosure.** On disclosure links $(k,\ell)\in\mathcal{R}_t$, the buyer receives $\hat{q}_{j\ell,t}$ directly, which reduces $\mathcal{V}_{k\ell,t}$. Optimal attention reallocates away from disclosure links toward non-disclosure links, reinforcing the price-based learning channel precisely where disclosure is absent.

4. **Upstream vs. downstream asymmetry.** The marginal value of downstream attention is second order (by Proposition 1(ii)), while the marginal value of upstream attention is first order. Consequently, firms allocate more attention to upstream links than to downstream links of comparable size.

Implications for the IC and Proposition 2. With non-cooperative signals, the cooperative surplus coefficient c_W in the proof of Proposition 2 depends on $\sigma^2_{\text{up},k\ell}$: a more precise non-cooperative signal reduces $\gamma_{ij,t}$ and hence $\omega_{ij,t}$, shrinking the cooperative surplus. This adds an additional comparative static to the baseline result: own-side importance $\ell^{(k)}_{k\ell}$ (the share of buyer k's purchases from supplier ℓ) raises attention through the trade volume channel, which increases non-cooperative signal precision, reduces $\omega_{ij,t}$, and tightens the IC. Hence higher own-side importance makes disclosure *less* likely—a sign reversal relative to partner-side importance that reflects the distinction between the sender's incentive to communicate and the receiver's incentive to acquire information independently.

H General equilibrium closure

The bilateral mechanism developed in Section 5.1 takes the nominal wage w_t as a common variable that enters every firm's marginal cost through the labor share α_i. This appendix closes the model in general equilibrium by specifying household preferences, monetary policy, and market clearing. The closure pins down w_t and validates that the bilateral results carry through to the full economy.

The general equilibrium structure follows La'O and Tahbaz-Salehi (2022), adapted to our information and disclosure environment. The key result is that monetary policy controls the nominal wage, so $\log w_t - \bar{w}$ is an exogenous aggregate variable from the perspective of individual firms. This justifies treating w_t as given in the bilateral analysis.

Household

A representative household has preferences over consumption C_t and labor L_t:

$$\mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\log C_t - \frac{L_t^{1+\varphi}}{1+\varphi} \right], \tag{63}$$

where $\beta \in (0,1)$ is the household's discount factor and $\varphi \geq 0$ is the inverse Frisch elasticity of labor supply. The household supplies labor to all firms at the competitive nominal wage w_t and consumes a final good C_t at nominal price P_t.

The consumption aggregate C_t is a Cobb–Douglas bundle over industry-level output:

$$C_t = \prod_{i=1}^{n} C_{i,t}^{\theta_i}, \qquad \sum_{i=1}^{n} \theta_i = 1, \tag{64}$$

where θ_i is the consumption weight of industry i and $C_{i,t}$ is a CES aggregate over firm-level varieties within industry i.[42]

The household's intratemporal optimality condition equates the marginal rate of substitution to the real wage:

$$L_t^{\varphi} = \frac{w_t}{P_t}. \tag{65}$$

[42]The within-industry CES aggregator for the household's consumption problem is distinct from the bilateral demand structure derived in Appendix D. The household aggregates over final goods; bilateral demand between firms is governed by switching costs. The two demand systems coincide at the industry level (Cobb–Douglas expenditure shares) but differ at the firm level.

Monetary policy

Monetary policy sets the nominal wage w_t directly. This is a standard modeling device in the dispersed-information literature.[43] The policy rule is

$$\log w_t = \bar{w} + \xi_t, \tag{66}$$

where $\bar{w}$ is the steady-state log wage and ξ_t is a monetary shock, which may be correlated with the vector of industry-level shocks $\{e_{i,t}\}_{i=1}^{n}$. The monetary shock ξ_t is *observed by all firms*: it is an aggregate variable that does not create additional dispersed information.

Under (66), the common component $q_{i,t}$ defined in (10) becomes

$$q_{i,t} = \alpha_i \, \xi_t + e_{i,t} + \sum_{j=1}^{n} a_{ij} \, q_{j,t}, \tag{67}$$

which has the closed form $q_{i,t} = \sum_j h_{ij}(\alpha_j \xi_t + e_{j,t})$, where h_{ij} is the (i,j) element of the Leontief inverse $H = (I - A)^{-1}$. Because ξ_t is common knowledge, the dispersed-information problem concerns only the industry-specific shocks $\{e_{i,t}\}$: firm k's uncertainty about $q_{i,t}$ reduces to uncertainty about the weighted sum $\sum_j h_{ij} e_{j,t}$.

Production technology

Firm k in industry i produces output using a Cobb–Douglas technology:

$$y_{ik,t} = z_{ik,t} \, l_{ik,t}^{\alpha_i} \prod_{j=1}^{n} m_{ijk,t}^{a_{ij}}, \tag{68}$$

where $l_{ik,t}$ is labor, $m_{ijk,t}$ is the intermediate input from industry j, and $z_{ik,t}$ is firm-specific productivity.

Market clearing

Labor market. The labor market clears at w_t:

$$L_t = \sum_{i=1}^{n} \int_0^1 l_{ik,t} \, dk. \tag{69}$$

From the Cobb–Douglas technology (68), firm k's labor demand is $l_{ik,t} = \alpha_i \, mc_{ik,t} \, y_{ik,t} / w_t$. Aggregating across firms and using the law of large numbers to eliminate idiosyncratic components:

$$L_t = \sum_{i=1}^{n} \alpha_i \, \frac{\bar{mc}_{i,t} \, Y_{i,t}}{w_t},$$

where $Y_{i,t} \equiv \int_0^1 y_{ik,t} \, dk$ is industry-level output and $\bar{mc}_{i,t}$ is the industry-average marginal cost.

[43] See La'O and Tahbaz-Salehi (2022) (Section III.A) and Angeletos and La' O (2020). The assumption that monetary policy controls the nominal wage rather than, say, the money supply or a nominal interest rate is without loss of generality in this class of models: any policy instrument that pins down a single nominal aggregate is equivalent, since the model has no nominal rigidities beyond dispersed information about real shocks.

Goods market. Each industry's output is allocated between intermediate use and final consumption:

$$Y_{i,t} = C_{i,t} + \sum_{j=1}^{n} \int_0^1 x_{ij,k,t}\, dk. \tag{70}$$

Log-linear equilibrium

To first order around the symmetric steady state, the equilibrium is characterized by two systems: a *pricing system* and a *belief system*.

Pricing system. From (17), the industry-level log price is $\log P_{i,t} = \bar{P}_i + \bar{\hat{q}}_{i,t}$ (as derived in (73)). The aggregate price index is

$$\log P_t = \sum_{i=1}^{n} \theta_i \log P_{i,t} = \bar{P} + \sum_{i=1}^{n} \theta_i\, \bar{\hat{q}}_{i,t}. \tag{71}$$

Combining with the labor supply condition (65):

$$\varphi \log L_t = \log w_t - \log P_t = \xi_t + \bar{w} - \bar{P} - \sum_{i=1}^{n} \theta_i\, \bar{\hat{q}}_{i,t}, \tag{72}$$

which determines aggregate labor (and hence output) as a function of the monetary shock and the vector of average beliefs.

Belief system. The average beliefs $\bar{\hat{q}}_{i,t}$ are determined by (54), which depends on the disclosure set $\mathcal{R}_t$ and the weight matrices $(B_t^{\text{exog}}, B_t^{\text{endo}})$. The belief system is self-contained: it depends on the true common components $\mathbf{q}_t$ and the signal structure, but not on quantities or the aggregate price level.[44]

The general equilibrium closure therefore does not affect the bilateral mechanism: the pricing and belief systems are determined by the signal structure and the disclosure set $\mathcal{R}_t$ independently of quantities and household preferences. The household side determines the *real allocation*—how belief accuracy translates into output and consumption—but not the *information structure* that governs disclosure and markups. This block-recursive property justifies the partial equilibrium approach of Sections 5.1–5.3.

Welfare

We derive the welfare cost of information frictions and markup distortions from first principles, following La'O and Tahbaz-Salehi (2022) and extending their formula to account for the markup wedge generated by price-based learning.

[44]This block-recursive structure is a consequence of the Cobb–Douglas technology and the log-linear approximation: quantities enter marginal costs only through the wage w_t, which is pinned down by monetary policy. In models with non-unit elasticity of substitution between inputs, quantities would feed back into marginal costs and the pricing and belief systems would be jointly determined.

Efficient benchmark. The efficient benchmark is the *constrained efficient allocation*: all firms observe $q_{i,t}$ perfectly ($\bar{\hat{q}}_{i,t} = q_{i,t}$ for all i) and every link discloses ($\mathcal{R}_t = \mathcal{R}^{\text{full}}$), so that every markup equals μ_i^* and there is no additional wedge from price-based learning. Denote the corresponding log real wage by $x^{**} \equiv \log w - \log P^{**}$, where P^{**} is the price level under full information and full disclosure at the symmetric steady state.

Under monopolistic competition with profits rebated to the household, $P^{**}C^{**} = \mu^* \, w \, L^{**}$ at the industry level (using the Cobb–Douglas factor-share condition), so $U'(x^{**}) \neq 0$: the markup wedge drives a gap between the marginal utility of labor income and the marginal disutility of work. The welfare cost of the equilibrium is measured relative to x^{**} and has both a *first-order* component from the additional markup distortion caused by non-disclosure, and *second-order* components from information frictions.

Step 1: Industry and aggregate price levels. The household's CES price index for industry i satisfies, to second order:

$$\log P_{i,t} = \overline{\log p_{ik,t}} - \frac{\varepsilon_i - 1}{2} \operatorname{Var}_k(\log p_{ik,t}) \,, \tag{73}$$

where the over-bar denotes the cross-firm mean within industry i. The mean log price is $\overline{\log p_{ik,t}} = \bar{P}_i + m_i + \bar{\hat{q}}_{i,t}$, with

$$m_i \equiv \sum_{j=1}^{n} a_{ij}(1 - r_t^{ji}) \, \delta_{ij,t}, \qquad \delta_{ij,t} \equiv \log \mu_{ij,t}^D - \log \mu_i^* = \log \frac{\varepsilon_i(1-\omega_{ij,t})}{\varepsilon_i(1-\omega_{ij,t}) - 1} - \log \frac{\varepsilon_i}{\varepsilon_i - 1}, \tag{74}$$

the exact log markup wedge on non-disclosure link (i, j). Since $\omega_{ij,t}$ can be large (the calibration yields $\omega > 0.50$ on some non-disclosure links), we use the exact expression (21) rather than the first-order approximation $\varepsilon_i \omega_{ij}/(\varepsilon_i - 1)^2$; the two coincide to first order in ω but differ substantially when ω is close to $(\varepsilon_i - 1)/\varepsilon_i$. Under the efficient benchmark, $m_i = 0$ and $\bar{\hat{q}}_{i,t} = q_{i,t}$, so $\log P_{i,t}^{**} = \bar{P}_i + q_{i,t} - \frac{\varepsilon_i - 1}{2} V_i^*$, where $V_i^* \equiv \operatorname{Var}_k^*(\log p_{ik,t})$ is idiosyncratic cost dispersion present even under full information. The aggregate price-level deviation is therefore

$$\log P_t - \log P_t^{**} = \underbrace{\bar{m}}_{\text{markup level}} - \underbrace{\tilde{e}_t}_{\text{belief error}} \underbrace{- \frac{1}{2} \sum_{i=1}^{n} \theta_i (\varepsilon_i - 1) \Delta V_{i,t}}_{\text{dispersion change}}, \tag{75}$$

where $\bar{m} \equiv \sum_i \theta_i m_i$, $\tilde{e}_t \equiv \sum_i \theta_i (q_{i,t} - \bar{\hat{q}}_{i,t})$, and $\Delta V_{i,t} \equiv \operatorname{Var}_k(\log p_{ik,t}) - V_i^* \geq 0$ is the excess within-industry price dispersion due to belief heterogeneity.

Step 2: Within-industry excess price dispersion. Firm k in industry i sets its price based on its own private signal $s_{ik,t} = q_{i,t} + \nu_{ik,t}$ only. The Kalman gain $\alpha_i = \sigma_{q,i}^2/(\sigma_{q,i}^2 + \sigma_{\nu,i}^2)$ gives the industry-average log price as $\overline{\log p_{ik,t}} = \bar{P}_i + m_i + \alpha_i q_{i,t}$, and each firm's deviation from the industry average is:

$$\log p_{ik,t} - \overline{\log p_{ik,t}} = \alpha_i \, \nu_{ik,t} + (\hat{\zeta}_{ik,t} - \bar{\hat{\zeta}}_{i,t}), \tag{76}$$

where $\nu_{ik,t}$ is firm k's private-signal noise and the $\hat{\zeta}$ terms capture idiosyncratic cost components (orthogonal to the common component q_i and present even under full information). The excess within-industry price dispersion due to imperfect information is therefore:

$$\Delta V_{i,t} = \alpha_i^2 \, \sigma_{\nu,i}^2. \tag{77}$$

This expression depends only on the seller's own-signal precision and is *constant across disclosure regimes*: network learning (price signals, disclosed beliefs) affects the buyer's demand but not the seller's price. Equation (77) replaces the full belief-system expression that would apply if prices depended on bilateral learning; we use the own-signal expression throughout the quantitative evaluation.

Note that the industry-average belief about q_i that enters the aggregate price level in (75) is $\bar{\hat{q}}_{i,t} = \alpha_i q_{i,t}$ (the own-signal posterior mean), so the aggregate belief error $\tilde{e}_t = \sum_i \theta_i(1-\alpha_i)q_{i,t}$ depends only on α_i and is likewise constant across regimes.

Step 3: Welfare expansion. Household welfare depends on the log real wage $x_t = \log w_t - \log P_t$ through (63) and (65). With Cobb–Douglas production and profits recycled to the household, $P_t C_t = \mu^* w_t L_t$ at the industry level, giving $\log C_t = \log \mu^* + x_t(1+\varphi)/\varphi + \text{const}$, so that:

$$U'(x) = \frac{1+\varphi}{\varphi} - \frac{1}{\varphi}\, e^{x(1+\varphi)/\varphi}, \qquad U''(x) = -\frac{(1+\varphi)^2}{\varphi^2}\, e^{x(1+\varphi)/\varphi}. \tag{78}$$

To second order around x^{**}:

$$\mathcal{L} \equiv \mathbf{E}[U(x^{**}) - U(x_t)] \approx -U'(x^{**})\,\mathbf{E}[x_t - x^{**}] - \tfrac{1}{2}U''(x^{**})\,\mathbf{E}[(x_t - x^{**})^2]. \tag{79}$$

The first-order term does not vanish because $U'(x^{**}) \neq 0$ (monopoly markup at the benchmark) *and* $\mathbf{E}[x_t - x^{**}] \neq 0$ (the markup wedge $\bar{m}$ deterministically depresses the real wage). Specifically, from (75):

$$\mathbf{E}[x_t - x^{**}] = -\mathbf{E}[\log P_t - \log P^{**}] = -\bar{m} + O(\omega^2), \tag{80}$$

since $\mathbf{E}[\tilde{e}_t] = 0$ (Gaussian belief errors are mean-zero) and the $\Delta V_{i,t}$ term is $O(\omega^2)$. Therefore the first-order term is $U'(x^{**}) \cdot \bar{m}$, which is $O(\omega)$ since $\bar{m} \propto \omega_{ij}$. For the second-order term, using $-U''(x^{**}) = (1+\varphi)^2 e^{x^{**}(1+\varphi)/\varphi}/\varphi^2$:

$$\mathbf{E}[(x_t - x^{**})^2] = \mathrm{Var}(\tilde{e}_t) + O(\omega^3), \tag{81}$$

where $\bar{m}^2$ is absorbed into $O(\omega^2)$ relative to the first-order term, and $\bar{m} \cdot \mathbf{E}[\tilde{e}_t] = 0$.

Step 4: Complete welfare formula.

$$\mathcal{L} = \underbrace{U'(x^{**})\,\bar{m}}_{\substack{\text{T1: markup cost of non-disclosure} \\ \text{(first order in } \omega\text{)}}} + \underbrace{\frac{(1+\varphi)^2}{2\varphi^2}\, e^{x^{**}(1+\varphi)/\varphi}\ \mathrm{Var}(\tilde{e}_t)}_{\substack{\text{T2: aggregate belief error} \\ \text{(La'O first term, second order)}}} + \underbrace{\frac{1}{2}\sum_{i=1}^{n} \theta_i(\varepsilon_i - 1)\,\mathbf{E}[\Delta V_{i,t}]}_{\substack{\text{T3: within-industry belief dispersion} \\ \text{(La'O second term, second order)}}}, \tag{82}$$

where $\tilde{e}_t = \sum_i \theta_i(q_{i,t} - \bar{\hat{q}}_{i,t})$, $\bar{m} = \sum_i \theta_i m_i$, $\Delta V_{i,t}$ is given by (77), and $U'(x^{**})$ is evaluated at the constrained-efficient steady state. We normalize $L^{**} = 1$ (a free choice of labor units), which gives $x^{**} = 0$ and $e^{x^{**}(1+\varphi)/\varphi} = 1$. With $\varphi = \frac{1}{2}$ (Frisch elasticity of 2) following La'O and Tahbaz-Salehi (2022), $U'(x^{**}) = 1/\varphi = 2$ and $\frac{(1+\varphi)^2}{2\varphi^2} = \frac{9}{2}$. These are the numerical coefficients used in the quantitative evaluation.[45]

[45]The formula captures welfare costs arising from seller-side belief distortions—the markup wedge on non-disclosure links (T1) and the seller pricing errors induced by imperfect information (T2, T3)—but abstracts from a fourth channel: the welfare cost of buyer switching decision errors. Buyers decide whether to switch to

Interpreting the three terms. *T1* is the *first-order markup cost of non-disclosure*. It is proportional to $\bar{m} = \sum_i \theta_i m_i$, the consumption-weighted average markup wedge induced by price-based learning on non-disclosure links. T1 is $O(\omega)$: it is the dominant welfare cost when ω_{ij} is not small, and it vanishes under full disclosure. Its coefficient $U'(x^{**}) > 0$ reflects the gap between the marginal utility of consumption and the marginal disutility of labor at the constrained-efficient steady state, which is positive precisely because the markup $\mu^* > 1$ keeps output below the social optimum.

T1 is related to but distinct from the bilateral surplus formula $\mathcal{L}_\mu = \sum_{(k,\ell)\notin\mathcal{R}_t} \Delta W_{k\ell,t}$ from (23). Both are first order in ω_{ij} and measure the markup distortion from non-disclosure; they differ in their network aggregation weights. $\mathcal{L}_\mu$ uses Domar weights $\lambda_i(1-\alpha_i)a_{ij}$, while T1 uses consumption weights $\theta_i a_{ij}$ through $\bar{m}$. The ratio $\text{T1}/\mathcal{L}_\mu = U'(x^{**}) \cdot \sum_i \theta_i m_i / \sum_{(i,j)} \lambda_i(1-\alpha_i)a_{ij}|W'_{ij}|\delta_{ij}$ is a model-specific constant determined by the steady-state network aggregation. We compute T1 directly from the GE formula (82) for Table 12, and report $\mathcal{L}_\mu$ as a supplementary quantity in Section 5.3.

T2 is the La'O aggregate information-loss term: stochastic deviations of industry-average prices from the truth generate aggregate price-level fluctuations that distort labor supply. Since prices depend only on the seller's own signal, $\tilde{e}_t = \sum_i \theta_i(1-\alpha_i)q_{i,t}$ and $\text{Var}(\tilde{e}_t)$ depends only on α_i and Σ_q. T2 is therefore *constant across disclosure regimes*: network learning affects demand-side beliefs but not prices, so the aggregate pricing error is invariant to the information regime.

T3 is the La'O within-industry dispersion term: belief heterogeneity across firms within the same industry generates within-industry price dispersion, which reduces real consumption through the CES aggregator. From (77), $\Delta V_{i,t} = \alpha_i^2 \sigma_{\nu,i}^2$ depends only on the seller's own-signal precision. Like T2, T3 is *constant across disclosure regimes*. It vanishes under full information (all firms observe q_i perfectly) but does not respond to changes in the bilateral learning structure.

---

the wholesale option based on $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ rather than the true $Q_{j,t}$; when beliefs are incorrect, some buyers switch who should not (paying iceberg waste $(\tau-1)Q_{j,t}$ unnecessarily) and some stay who should switch (paying a distorted markup). The welfare cost of this misallocation is proportional to the density of buyers at the switching threshold multiplied by the variance of buyer belief errors about $Q_{j,t}$, and is therefore second order in belief errors. In the calibrated model, buyer belief errors are moderate (Kalman gain $\alpha \approx 0.48$) and the Pareto switching cost distribution with $\varepsilon = 6$ places little mass near the marginal buyer, so this channel is quantitatively negligible relative to T1. The approximation follows the standard practice in this class of models of treating switching probabilities as constant to first order (Section 5.1), and implies that T2 and T3 are slightly understated; the overall welfare cost (82) is therefore a conservative lower bound on the full second-order cost.

Online Appendix

(not for publication)

I Variance decomposition and counterfactual accounting

This appendix quantifies the economic magnitude of partner-based learning using two complementary exercises. First, we report incremental explanatory power from partner beliefs in the baseline specification. Second, we construct a simple counterfactual that asks how much expectations would change if all currently untreated partners were instead treated.

Throughout, we restrict attention to untreated buyers ($T_i = 0$) and use uniform within-firm weights.

I.1 Variance decomposition

We measure the incremental R^2 contribution of partner beliefs by comparing the R^2 from the baseline specification (without partner variables) to the R^2 from the augmented specification that adds partner beliefs interacted with treatment status. Because the outcome signal differs in scale across directions—supplier prices versus customer purchase growth—we focus on the belief channel, which uses the same type of regressor (partner inflation expectations) on both sides and therefore permits a clean cross-direction comparison. Within each direction, the specification includes firm fixed effects, sector-by-quarter, size-by-quarter, and area-by-quarter fixed effects, along with lagged own expectations and a horizon indicator.

Table I.1 reports the results for uniform weights. Adding partner beliefs increases the R^2 by 0.18 percentage points when both directions are pooled. The explanatory contribution of partner beliefs is concentrated in the customer direction: beliefs from customers account for a 0.14 pp increase in R^2, compared to 0.05 pp from supplier beliefs. These magnitudes are modest in absolute terms but consistent with the counterfactual analysis below, which provides a complementary assessment of economic significance using estimated coefficients from the full specification.

Table I.1: Incremental R^2 from partner beliefs

	Suppliers only	Customers only	Both directions
Partner beliefs	0.0005	0.0014	0.0018

Notes: Entries report the increase in R^2 (in absolute terms) relative to the baseline specification without partner variables. Regressions are run separately within the supplier and customer samples and use uniform within-firm weights. "Both directions" sums the direction-specific ΔR^2 values. Sample restricted to untreated buyers.

I.2 Counterfactual: all untreated partners treated

We next construct a linear counterfactual using the estimated pooled-direction coefficients. For each untreated firm i, we compute the predicted change in its expectation if all currently untreated partners were instead treated. Formally, the counterfactual effect equals the sum over untreated partners j of the estimated treated-partner slope multiplied by the realized partner signal and the corresponding link weight.

Table I.2 reports the average predicted change in expectations among firms with at least one untreated partner. The mean counterfactual change is 0.041 percentage points. Decomposing by direction, 0.019 pp arise from suppliers and 0.022 pp from customers. Decomposing by

channel, the dominant contribution comes from partner beliefs, while outcome signals account for a smaller share. In particular, supplier beliefs contribute 0.014 pp, customer beliefs 0.021 pp, supplier prices 0.005 pp, and customer purchase signals 0.001 pp.

These magnitudes are modest in absolute terms but economically meaningful relative to the typical quarterly change in expectations in our sample. They reinforce the conclusion that information transmitted through partner beliefs constitutes the primary channel of network-based learning.

Table I.2: Counterfactual change in expectations if all untreated partners were treated

	Mean predicted change (pp)
Total effect	0.0408
Supplier component	0.0189
Customer component	0.0219
Supplier beliefs	0.0143
Supplier prices	0.0046
Customer beliefs	0.0214
Customer outcomes	0.0005

Notes: Counterfactual constructed using estimated pooled-direction coefficients from the full specification. The sample includes untreated firms with at least one untreated partner. Reported values are means of the predicted change in expectations (in percentage points).

J The SIGE survey: design and variable construction

This appendix provides additional details on the *Survey on Inflation and Growth Expectations* (SIGE) and describes the construction of the main variables used in the empirical analysis. A complete description of the survey, including the full questionnaire, is provided in Coibion et al. (2020) and their Online Appendix.

Sampling frame. The SIGE is a quarterly business survey conducted by the Bank of Italy since December 1999. The reference universe consists of firms with administrative headquarters in Italy employing 50 or more workers, operating in industry (excluding construction until 2013:Q1), nonfinancial private services, and—since 2013:Q1—construction.[46] The sample is stratified along three dimensions: sector of economic activity (industry, nonfinancial private services, construction), geographic area (northwest, northeast, center, south and islands), and firm size (50–199, 200–999, 1,000 and over employees). In recent waves, each round includes approximately 1,000–1,300 firms. The sampling frame is drawn from Bureau Van Dijk's Aida database and updated approximately every five years. Sampling weights are provided to ensure that the distribution of firms by employment in the sample represents the population distribution.

[46]The following sectors are excluded: financial intermediaries and insurance companies, general government, educational and health care services, and other community, social, and personal services.

Fieldwork and mode of administration. The survey is administered by a specialist firm during the first three weeks of March, June, September, and December. Questionnaires are distributed to company managers. Approximately 90% of responses are collected through computer-assisted web interviews (CAWI), with the remaining 10% collected by computer-assisted telephone interviews (CATI). The average response rate is approximately 45%.

Survey content. The questionnaire covers inflation expectations, assessments of the general economic situation, own-product prices and demand, investment plans, and employment. Most responses are qualitative, with the key exceptions being own-product price changes (past and expected) and inflation expectations, which are elicited as quantitative point forecasts. A representative questionnaire is reproduced in Coibion et al. (2020), Online Appendix Figure 1.

Variables used in our analysis. We use the following SIGE variables:

- *Inflation expectations.* Quantitative point forecasts of Italian CPI inflation at the 6-month and 12-month horizons (survey questions B1a and B1b). These are the dependent variables in the network learning regressions.
- *Own price growth.* The average year-on-year percentage change in the prices of goods billed by the firm (*"Variazione percentuale media annua dei prezzi dei beni da Voi fatturati"*; survey question D1). This variable serves as the outcome-based signal in the upstream learning regressions and is the basis for the price-based forecast-accuracy diagnostics in Section 3.1.
- *Firm characteristics.* Employment (number of employees), sectoral classification, and geographic area, used to construct the interacted fixed effects $\alpha_{\text{sector}\times t}$, $\alpha_{\text{size}\times t}$, and $\alpha_{\text{area}\times t}$ in the estimation.

Information treatment. The randomized information treatment was introduced in 2012. In each wave, treated firms are shown the following information immediately before reporting their inflation expectations: the 12-month change in the harmonized index of consumer prices (HICP) for Italy and for the euro area in the most recent available month (survey question B1a, treated version). Untreated firms proceed directly to the expectations questions without receiving this information. Approximately two-thirds of participating firms are assigned to treatment.

Treatment assignment is persistent within re-randomization windows. The initial randomization took place in 2012, with treatment status remaining fixed until a re-randomization before 2019. In our sample period (2019–2023), treatment status is therefore predetermined and time-invariant. Coibion et al. (2020), Online Appendix Table 2, show that treatment assignment is uncorrelated with firm size, sector, geographic area, export intensity, and pricing behavior.

Survey weights. Throughout the analysis, we use the survey's analytic weights when available (otherwise set to 1). In the network regressions, these firm–quarter weights are distributed across a firm's observed trading partners within each quarter, ensuring that the total weight attached to firm i in quarter t equals its survey weight regardless of the number of observed links (see Section 3.2 for details).

K Administrative transaction data and network construction

This appendix provides additional details on the administrative transaction data used to construct the buyer–supplier networks and describes the matching procedure linking these records to the SIGE survey.

Data source. The transaction data originate from Italy's mandatory electronic invoicing system (*fatturazione elettronica*), administered by the Italian Revenue Agency (*Agenzia delle Entrate*). Since 1 January 2019, all business-to-business transactions in goods and services between VAT-registered entities resident, established, or registered in Italy must be reported electronically. Each invoice record includes the tax identifiers of the buyer and seller, the taxable amount (*imponibile*), the VAT amount, the invoice date, and a sectoral classification of the issuing firm. We use the taxable amount as our measure of transaction value throughout.

Aggregation and timing. The raw data are at the invoice level. We aggregate all invoices between a given buyer–seller pair within each calendar quarter to obtain quarterly bilateral transaction values, aligning the network data with the quarterly frequency of the SIGE survey. The fieldwork periods of SIGE (first three weeks of March, June, September, and December) fall within the same quarter as the transaction flows they are matched to.

Matching to SIGE. We match SIGE firms to the electronic invoicing records using the firm's tax identifier (*codice fiscale*). A firm–partner–quarter observation (i, j, t) enters the analysis sample if both i and j appear in SIGE in quarter t. Because SIGE covers roughly 1,000–1,300 firms per wave while the invoicing universe contains millions of VAT-registered entities, only a small share of each firm's trading partners are observed in the survey. At the median firm–quarter, surveyed suppliers account for approximately 2.6% of all upstream trading relationships.

Directional structure. Each matched pair (i, j, t) generates two directional observations: one in which j is a supplier of i (upstream, $d = \text{up}$) and one in which j is a customer of i (downstream, $d = \text{dn}$). The outcome signal differs by direction: for upstream links, the signal is supplier j's year-on-year own price growth $g^p_{j,t}$ as reported in SIGE; for downstream links, it is customer j's year-on-year purchase growth $g^r_{j,t}$, constructed from the administrative transaction values.

Network variables. For each firm–partner–quarter–direction observation, we construct:

- *Transaction value.* The total taxable amount (*imponibile*) aggregated over the quarter along the directed link (i, j, t, d).
- *Partner-side share (value-based).* The share of partner j's total transaction value cumulated over the sample accounted for by firm i: $s^{(j)}_{ij} \equiv \bar{v}_{ij} / \sum_{i'} \bar{v}_{i'j}$, where $\bar{v}_{ij}$ is the bilateral taxable amount summed across all quarters.
- *Partner-side share (persistence-based).* The number of quarters in which dyad (i, j) is observed, expressed as a share of partner j's total link-quarters across all partners.
- *Partner breadth.* The number of distinct trading partners of firm j in quarter t, measured direction-specifically: for supplier dyads, j's number of upstream suppliers; for customer dyads, j's number of downstream customers.

- *Link weights.* Within-firm weights that allocate each firm–quarter survey weight $w_{i,t}$ across observed partners. Under uniform weighting, each link receives equal weight $\omega_{ij,t} = w_{i,t}/|\mathcal{N}_{i,t}|$; under transaction-value weighting, weights are proportional to bilateral trade value; under breadth weighting, weights reflect the partner's network centrality.

Sample restrictions. The baseline estimation sample restricts attention to fourth-quarter observations to provide a consistent annual cross-section and to mitigate seasonality. All continuous variables entering the regressions (inflation expectations, own price growth, purchase growth) are winsorized at the 1st and 99th percentiles using analysis weights.

L Separate-horizon/direction estimates

This appendix reports horizon-specific versions of the learning-from-transactions specifications. We estimate the supplier-side and customer-side regressions separately at $h \in \{6, 12\}$, keeping the same fixed effects, dyad controls, and inference procedures as in the pooled-horizon analysis. We report results under both *global* and *quarter-level* link-weight normalizations.

Setup and variables. We use dyad–quarter observations (i, j, t) drawn from the administrative transaction network. The dependent variable is firm i's h-quarter-ahead inflation expectation in quarter t, $\pi_{i,t}^{e,h}$. For each partner j, we consider two partner-generated signals observed in the same quarter: (i) an outcome signal and (ii) the partner's inflation expectation (belief signal).

In the upstream analysis (suppliers), the outcome signal is the supplier's year-on-year own-price growth, $g_{j,t}^{p}$. In the downstream analysis (customers), the outcome signal is the customer's year-on-year purchase growth, $g_{j,t}^{r}$, constructed from winsorized taxable amount growth and expressed in percentage points. In both directions, the belief signal is the partner's h-ahead inflation expectation, $\pi_{j,t}^{e,h}$.

Weights and column structure. Regressions use analytic weights that allocate the firm–quarter survey weight across observed links within direction (suppliers or customers). We consider two alternative normalization schemes. Under *global* normalization, link weights are constructed using the full network-linked sample window and then held fixed across time. Under *quarter-level* normalization, weights are renormalized separately within each firm–quarter, so that link importance reflects contemporaneous transaction patterns.

Within each normalization scheme, we report four alternative within-direction weighting rules: baseline partner-exposure (breadth or centrality) weights (brd), uniform link weights (unif), transaction-value weights based on taxable amount shares (val), and transaction-value weights tilted by the buyer's share of the partner's total sales in that quarter (shr).

Each table contains nine columns. Column (1) reports the transaction-only specification. Columns (2)–(5) add the partner belief signal (Untreated only). Columns (6)–(9) further include interactions with the focal firm's treatment status T_i, allowing treated and untreated firms to load differently on treated and untreated partners' signals.

Interpretation. Because the pooled-horizon analysis imposes common slopes across horizons, while the specifications reported here allow coefficients to vary freely by horizon, the estimates are not directly comparable in magnitude. Instead, the horizon-specific tables are intended as a

robustness exercise that assesses whether the qualitative structure of information transmission is stable across horizons and weighting schemes.

Main patterns. The horizon-specific estimates display several regularities. First, in both the supplier-side and customer-side regressions, the coefficient on $T_j \times \pi_{j,t}^{e,h}$ is typically positive across weighting schemes, indicating stronger loading on treated partners' beliefs. Second, coefficients on outcome-based signals are generally smaller in magnitude and more sensitive to the choice of link-weight normalization. Third, estimates at $h = 12$ are often attenuated relative to $h = 6$, although the direction of effects is unchanged in most specifications.

Taken together, these patterns indicate that the qualitative structure of spillovers documented in the pooled-horizon analysis—namely, the prominence of belief-based transmission relative to outcome-based transmission—also appears in the horizon-specific regressions.

Horizon-specific estimates. The tables below report horizon-specific spillover estimates separately for suppliers and customers. For each horizon, we present results under global link-weight normalizations and under quarter-level normalizations. These estimates are not intended to replace the pooled specification in the main text, but to document the robustness of the underlying spillover patterns.

L.1 Robustness: $h = 12$ evidence on relationship size

This appendix examines the robustness of the uniform-weight heterogeneity results in Section 4 when attention is restricted to a single forecast horizon. Specifically, we re-estimate the relationship-size specifications using only the 12-quarters-ahead inflation expectation, $\pi_{i,t}^{e,12}$, rather than pooling across horizons.

Design. The empirical design mirrors the pooled analysis as closely as possible. The dependent variable is firm i's 12-quarter-ahead inflation expectation reported in quarter t, denoted $y_{ij,t} \equiv \pi_{i,t}^{e,12}$. There is no horizon stacking and hence no horizon fixed effects. The lag structure remains horizon-specific and we include the appropriate lagged outcome $y_{ij,t-1} = \pi_{i,t-1}^{e,12}$. All specifications focus on supplier relationships (links in which firm j supplies firm i), restrict attention to untreated buyers ($T_i = 0$), use uniform within-buyer–quarter weights, and include the same fixed effects and controls as in the pooled benchmark.

Linear interaction estimates. Appendix Table L.9 reports linear interaction estimates for the $h = 12$ outcome, separately for supplier-side and buyer-side size definitions and for price-only versus horse-race specifications. The qualitative patterns closely mirror those in the pooled analysis. In particular, belief-based learning from treated suppliers exhibits clear heterogeneity with respect to relationship size when size is defined on the supplier side. The interaction between treated suppliers' beliefs and supplier-side size remains economically meaningful and statistically significant, indicating that untreated buyers load more heavily on the beliefs of treated suppliers precisely in relationships that are more important for the supplier.

By contrast, heterogeneity in the price channel is weaker. While the average treated-supplier price effect remains positive, its dependence on relationship size is modest and less stable across size definitions. These findings reinforce the interpretation that belief transmission is governed

Table L.1: Learning from suppliers ($h = 6$), global link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner price growth $g^p_{j,t}$	-0.014^{***}	-0.008	-0.011	-0.008	-0.002	-0.008	-0.012	-0.009	-0.002
	(0.005)	(0.006)	(0.008)	(0.006)	(0.008)	(0.006)	(0.008)	(0.006)	(0.008)
$T_j \times g^p_{j,t}$	0.001	0.006	0.015^{**}	0.005	-0.000	0.005	0.010	0.003	-0.001
	(0.007)	(0.009)	(0.006)	(0.007)	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.185^{***}	0.167^{***}	0.124^{*}	0.109^{*}	0.112^{*}	0.206^{***}	0.207^{***}	0.206^{***}	0.206^{***}
	(0.063)	(0.058)	(0.069)	(0.061)	(0.062)	(0.047)	(0.048)	(0.047)	(0.047)
$T_i \times g^p_{j,t}$						0.002	0.018^{*}	0.011	-0.007
						(0.007)	(0.009)	(0.007)	(0.009)
$T_i \times T_j \times g^p_{j,t}$						0.001	-0.016^{*}	0.003	0.001
						(0.011)	(0.009)	(0.011)	(0.010)
Partner belief $\pi^{e,h}_{j,t}$		-0.003	-0.002	0.000	0.018	-0.003	-0.004	0.000	0.019
		(0.019)	(0.034)	(0.029)	(0.027)	(0.019)	(0.034)	(0.028)	(0.027)
$T_j \times \pi^{e,h}_{j,t}$		0.058^{***}	0.059^{***}	0.111^{***}	0.068^{***}	0.058^{***}	0.066^{***}	0.106^{***}	0.065^{***}
		(0.021)	(0.019)	(0.036)	(0.022)	(0.021)	(0.021)	(0.031)	(0.022)
$T_i \times \pi^{e,h}_{j,t}$						0.001	-0.020	0.009	-0.022
						(0.030)	(0.044)	(0.034)	(0.030)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						-0.009	0.005	-0.014	0.002
						(0.029)	(0.030)	(0.031)	(0.023)
Observations	140,366	140,366	45,046	45,046	45,046	140,366	140,366	140,366	140,366
R^2	0.039	0.043	0.033	0.029	0.028	0.053	0.050	0.053	0.053

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on partner signals observed in the same quarter. For suppliers, $g^p_{j,t}$ is supplier j's year-on-year own price growth. Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j. Columns (2)–(5) add partner beliefs under alternative link-weight normalizations: breadth/centrality (brd), uniform (unif), transaction value (val), and partner sales share (shr). Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) include T_i interactions. All specifications include baseline controls and fixed effects and report Driscoll–Kraay standard errors (bandwidth 2); firm-level clustering is used when the variance matrix is not positive definite.

Table L.2: Learning from suppliers ($h = 12$), global link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner price growth $g^p_{j,t}$	−0.013**	−0.008	−0.008	−0.003	0.005	−0.008	−0.008	−0.003	0.005
	(0.006)	(0.007)	(0.007)	(0.005)	(0.008)	(0.007)	(0.007)	(0.005)	(0.008)
$T_j \times g^p_{j,t}$	0.004	0.009	0.015**	0.005	−0.001	0.009	0.013*	0.006	−0.001
	(0.009)	(0.010)	(0.006)	(0.007)	(0.008)	(0.010)	(0.007)	(0.008)	(0.008)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.222***	0.212***	0.181**	0.182***	0.179**	0.238***	0.234***	0.238***	0.239***
	(0.062)	(0.057)	(0.086)	(0.070)	(0.072)	(0.050)	(0.048)	(0.050)	(0.050)
$T_i \times g^p_{j,t}$						0.002	−0.002	0.004	−0.011
						(0.009)	(0.006)	(0.007)	(0.008)
$T_i \times T_j \times g^p_{j,t}$						−0.002	−0.004	0.002	0.008
						(0.009)	(0.012)	(0.009)	(0.009)
Partner belief $\pi^{e,h}_{j,t}$		0.016	0.010	0.012	0.025	0.016	0.010	0.012	0.030
		(0.017)	(0.025)	(0.040)	(0.033)	(0.017)	(0.026)	(0.033)	(0.032)
$T_j \times \pi^{e,h}_{j,t}$		0.027	0.033**	0.074**	0.039**	0.027	0.041**	0.069**	0.036*
		(0.019)	(0.015)	(0.033)	(0.017)	(0.019)	(0.018)	(0.027)	(0.019)
$T_i \times \pi^{e,h}_{j,t}$						0.018	0.020	0.031	−0.004
						(0.027)	(0.029)	(0.031)	(0.032)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.024	−0.035*	−0.038	0.001
						(0.023)	(0.021)	(0.024)	(0.016)
Observations	140366	140366	446?	44702	44702	140666	140666	140666	140666
R^2	0.052	0.054	0.033	0.037	0.036	0.051	0.049	0.052	0.052

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on partner signals observed in the same quarter. For suppliers, $g^p_{j,t}$ is supplier j's year-on-year own price growth as reported in the survey. Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j ("Trans. only"). Columns (2)–(5) add the partner belief signal (Untreated only) under four alternative within-firm link-weight normalizations: "brd" uses the baseline breadth/centrality weights within supplier links, "unif" assigns uniform weights across observed supplier links, "val" weights links by transaction value (taxable amount shares), and "shr" tilts transaction-value weights by the buyer's share of partner j's total sales in that quarter. Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) use the full sample and include T_i interactions (triple-interaction specifications). All specifications include the baseline controls and fixed effects used in the main text (buyer fixed effects and time fixed effects, plus the additional interacted fixed effects), use analytic weights, and report standard errors using the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

Table L.3: Learning from customers ($h = 6$), global link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner purchase growth $g^r_{j,t}$	−0.000	−0.001	−0.001	0.000	−0.000	−0.001	−0.001	0.000	−0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$T_j \times g^r_{j,t}$	0.001	0.001	0.001	0.001	0.002	0.001	0.001	0.000	0.001
	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.220***	0.215***	0.189***	0.204***	0.201***	0.238***	0.235***	0.239***	0.239***
	(0.051)	(0.048)	(0.053)	(0.059)	(0.057)	(0.039)	(0.039)	(0.039)	(0.039)
$T_i \times g^r_{j,t}$						0.002	0.001	0.001	0.000
						(0.003)	(0.002)	(0.003)	(0.002)
$T_i \times T_j \times g^r_{j,t}$						−0.003	−0.001	−0.003	−0.002
						(0.004)	(0.003)	(0.004)	(0.003)
Partner belief $\pi^{e,h}_{j,t}$		−0.007	−0.016	−0.030	−0.011	−0.007	−0.016	−0.034	−0.011
		(0.028)	(0.029)	(0.029)	(0.030)	(0.028)	(0.029)	(0.032)	(0.030)
$T_j \times \pi^{e,h}_{j,t}$		0.066**	0.059***	0.105***	0.066**	0.066**	0.058***	0.102***	0.061**
		(0.029)	(0.018)	(0.029)	(0.031)	(0.029)	(0.018)	(0.031)	(0.031)
$T_i \times \pi^{e,h}_{j,t}$						0.048	0.046**	0.077**	0.046
						(0.034)	(0.019)	(0.034)	(0.033)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.060	−0.053**	−0.065	−0.055
						(0.043)	(0.023)	(0.044)	(0.036)
Observations	108,187	108,187	38,452	38,452	38,452	108,187	108,187	108,187	108,187
R^2	0.070	0.071	0.066	0.055	0.055	0.070	0.070	0.071	0.070

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on partner signals observed in the same quarter. For customers, $g^r_{j,t}$ is customer j's year-on-year purchase growth (winsorized at the 1st and 99th percentiles and expressed in percentage points). Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j ("Trans. only"). Columns (2)–(5) add the partner belief signal (Untreated only) under four alternative within-firm link-weight normalizations: "brd" uses the baseline breadth/centrality weights within customer links, "unif" assigns uniform weights across observed customer links, "val" weights links by transaction value (taxable amount shares), and "shr" tilts transaction-value weights by the buyer's share of partner j's total sales in that quarter. Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) use the full sample and include T_i interactions (triple-interaction specifications). All specifications include the baseline controls and fixed effects used in the main text (buyer fixed effects and time fixed effects, plus the additional interacted fixed effects), use analytic weights, and report standard errors using the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

Table L.4: Learning from customers ($h = 12$), global link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner purchase growth $g^r_{j,t}$	0.001	0.001	0.000	0.001	0.001	0.001	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$T_j \times g^r_{j,t}$	−0.001	−0.001	0.002	−0.002	0.002	−0.001	0.001	−0.002	0.001
	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.225***	0.225***	0.207***	0.209***	0.205***	0.241***	0.238***	0.241***	0.240***
	(0.050)	(0.047)	(0.047)	(0.062)	(0.062)	(0.041)	(0.041)	(0.041)	(0.041)
$T_i \times g^r_{j,t}$						0.001	0.002	0.001	0.002
						(0.003)	(0.002)	(0.003)	(0.002)
$T_i \times T_j \times g^r_{j,t}$						−0.001	−0.004	−0.002	−0.002
						(0.004)	(0.003)	(0.004)	(0.003)
Partner belief $\pi^{e,h}_{j,t}$		−0.013	−0.012	−0.028	−0.011	−0.013	−0.012	−0.032	−0.011
		(0.018)	(0.018)	(0.026)	(0.028)	(0.018)	(0.018)	(0.026)	(0.028)
$T_j \times \pi^{e,h}_{j,t}$		0.065***	0.060***	0.105***	0.073**	0.065***	0.057***	0.102***	0.068**
		(0.017)	(0.017)	(0.029)	(0.031)	(0.017)	(0.017)	(0.031)	(0.031)
$T_i \times \pi^{e,h}_{j,t}$						0.071**	0.050**	0.090**	0.052
						(0.032)	(0.019)	(0.032)	(0.033)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.068	−0.060**	−0.070	−0.059
						(0.042)	(0.023)	(0.043)	(0.036)
Observations	108187	108187	38170	38170	38170	108187	108187	108187	108187
R^2	0.073	0.072	0.069	0.057	0.057	0.072	0.070	0.073	0.071

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on partner signals observed in the same quarter. For customers, $g^r_{j,t}$ is customer j's year-on-year purchase growth (winsorized at the 1st and 99th percentiles and expressed in percentage points). Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j ("Trans. only"). Columns (2)–(5) add the partner belief signal (Untreated only) under four alternative within-firm link-weight normalizations: "brd" uses the baseline breadth/centrality weights within customer links, "unif" assigns uniform weights across observed customer links, "val" weights links by transaction value (taxable amount shares), and "shr" tilts transaction-value weights by the buyer's share of partner j's total sales in that quarter. Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) use the full sample and include T_i interactions (triple-interaction specifications). All specifications include the baseline controls and fixed effects used in the main text (buyer fixed effects and time fixed effects, plus the additional interacted fixed effects), use analytic weights, and report standard errors using the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

Table L.5: Learning from suppliers ($h = 6$), quarter link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner price growth $g^p_{j,t}$	−0.014***	−0.008	−0.011	−0.008	−0.002	−0.008	−0.012	−0.009	−0.002
	(0.005)	(0.006)	(0.008)	(0.006)	(0.008)	(0.006)	(0.008)	(0.006)	(0.008)
$T_j \times g^p_{j,t}$	0.001	0.006	0.015**	0.005	−0.000	0.005	0.010	0.003	−0.001
	(0.007)	(0.009)	(0.006)	(0.007)	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.185***	0.167***	0.124*	0.109*	0.112*	0.206***	0.207***	0.206***	0.206***
	(0.063)	(0.058)	(0.069)	(0.061)	(0.062)	(0.047)	(0.048)	(0.047)	(0.047)
$T_i \times g^p_{j,t}$						0.002	0.018*	0.011	−0.007
						(0.007)	(0.009)	(0.007)	(0.009)
$T_i \times T_j \times g^p_{j,t}$						0.001	−0.016*	0.003	0.001
						(0.011)	(0.009)	(0.011)	(0.010)
Partner belief $\pi^{e,h}_{j,t}$		−0.003	−0.002	0.000	0.018	−0.003	−0.004	0.000	0.019
		(0.019)	(0.034)	(0.029)	(0.027)	(0.019)	(0.034)	(0.028)	(0.027)
$T_j \times \pi^{e,h}_{j,t}$		0.058***	0.059***	0.111***	0.068***	0.058***	0.066***	0.106***	0.065***
		(0.021)	(0.019)	(0.036)	(0.022)	(0.021)	(0.021)	(0.031)	(0.022)
$T_i \times \pi^{e,h}_{j,t}$						0.001	−0.020	0.009	−0.022
						(0.030)	(0.044)	(0.034)	(0.030)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.009	0.005	−0.014	0.002
						(0.029)	(0.030)	(0.031)	(0.023)
Observations	140,366	140,366	45,046	45,046	45,046	140,366	140,366	140,366	140,366
R^2	0.039	0.043	0.033	0.029	0.028	0.053	0.050	0.053	0.053

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on partner signals observed in the same quarter. For suppliers, $g^p_{j,t}$ is supplier j's year-on-year own price growth as reported in the survey. Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j ("Trans. only"). Columns (2)–(5) add the partner belief signal (Untreated only) under four alternative within-firm link-weight normalizations: "brd" uses the baseline breadth/centrality weights within supplier links, "unif" assigns uniform weights across observed supplier links, "val" weights links by transaction value (taxable amount shares), and "shr" tilts transaction-value weights by the buyer's share of partner j's total sales in that quarter. Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) use the full sample and include T_i interactions (triple-interaction specifications). All specifications include the baseline controls and fixed effects used in the main text (buyer fixed effects and time fixed effects, plus the additional interacted fixed effects), use analytic weights, and report standard errors using the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

Table L.6: Learning from suppliers ($h = 12$), quarter-normalized link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner price growth $g^p_{j,t}$	−0.009	−0.006	−0.010	−0.007	−0.001	−0.006	−0.009	−0.007	−0.001
	(0.006)	(0.007)	(0.008)	(0.006)	(0.008)	(0.007)	(0.008)	(0.006)	(0.008)
$T_j \times g^p_{j,t}$	0.004	0.007	0.013*	0.004	0.001	0.007	0.011	0.003	0.000
	(0.008)	(0.009)	(0.007)	(0.008)	(0.008)	(0.009)	(0.008)	(0.008)	(0.008)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.198***	0.184***	0.149**	0.141**	0.143**	0.226***	0.227***	0.226***	0.226***
	(0.061)	(0.058)	(0.067)	(0.061)	(0.062)	(0.046)	(0.046)	(0.046)	(0.046)
$T_i \times g^p_{j,t}$						0.004	0.016*	0.010	−0.006
						(0.007)	(0.009)	(0.007)	(0.009)
$T_i \times T_j \times g^p_{j,t}$						−0.003	−0.014*	0.002	0.001
						(0.010)	(0.009)	(0.010)	(0.010)
Partner belief $\pi^{e,h}_{j,t}$		−0.004	−0.003	0.001	0.017	−0.004	−0.005	0.001	0.018
		(0.020)	(0.033)	(0.029)	(0.027)	(0.020)	(0.033)	(0.028)	(0.027)
$T_j \times \pi^{e,h}_{j,t}$		0.054**	0.056***	0.104***	0.062***	0.054**	0.061***	0.099***	0.060***
		(0.022)	(0.020)	(0.035)	(0.022)	(0.022)	(0.022)	(0.031)	(0.022)
$T_i \times \pi^{e,h}_{j,t}$						−0.002	−0.018	0.008	−0.020
						(0.029)	(0.043)	(0.033)	(0.029)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.011	0.004	−0.012	0.001
						(0.028)	(0.029)	(0.030)	(0.022)
Observations	140,366	140,366	45,046	45,046	45,046	140,366	140,366	140,366	140,366
R^2	0.042	0.046	0.036	0.032	0.031	0.056	0.053	0.056	0.056

Table L.7: Learning from customers ($h = 6$), quarter link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner purchase growth $g^r_{j,t}$	−0.000	−0.001	−0.001	0.000	−0.000	−0.001	−0.001	0.000	−0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$T_j \times g^r_{j,t}$	0.001	0.001	0.001	0.001	0.002	0.001	0.001	0.000	0.001
	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.220***	0.215***	0.189***	0.204***	0.201***	0.238***	0.235***	0.239***	0.239***
	(0.051)	(0.048)	(0.053)	(0.059)	(0.057)	(0.039)	(0.039)	(0.039)	(0.039)
$T_i \times g^r_{j,t}$						0.002	0.001	0.001	0.000
						(0.003)	(0.002)	(0.003)	(0.002)
$T_i \times T_j \times g^r_{j,t}$						−0.003	−0.001	−0.003	−0.002
						(0.004)	(0.003)	(0.004)	(0.003)
Partner belief $\pi^{e,h}_{j,t}$		−0.007	−0.016	−0.030	−0.011	−0.007	−0.016	−0.034	−0.011
		(0.028)	(0.029)	(0.029)	(0.030)	(0.028)	(0.029)	(0.032)	(0.030)
$T_j \times \pi^{e,h}_{j,t}$		0.066**	0.059***	0.105***	0.066**	0.066**	0.058***	0.102***	0.061**
		(0.029)	(0.018)	(0.029)	(0.031)	(0.029)	(0.018)	(0.031)	(0.031)
$T_i \times \pi^{e,h}_{j,t}$						0.048	0.046**	0.077**	0.046
						(0.034)	(0.019)	(0.034)	(0.033)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.060	−0.053**	−0.065	−0.055
						(0.043)	(0.023)	(0.044)	(0.036)
Observations	108,187	108,187	38,452	38,452	38,452	108,187	108,187	108,187	108,187
R^2	0.070	0.071	0.066	0.055	0.055	0.070	0.070	0.071	0.070

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on partner signals observed in the same quarter. For customers, $g^r_{j,t}$ is customer j's year-on-year purchase growth (winsorized at the 1st and 99th percentiles and expressed in percentage points). Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j ("Trans. only"). Columns (2)–(5) add the partner belief signal (Untreated only) under four alternative within-firm link-weight normalizations: "brd" uses the baseline breadth/centrality weights within customer links, "unif" assigns uniform weights across observed customer links, "val" weights links by transaction value (taxable amount shares), and "shr" tilts transaction-value weights by the buyer's share of partner j's total sales in that quarter. Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) use the full sample and include T_i interactions (triple-interaction specifications). All specifications include the baseline controls and fixed effects used in the main text (buyer fixed effects and time fixed effects, plus the additional interacted fixed effects), use analytic weights, and report standard errors using the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

Table L.8: Learning from customers ($h = 12$), quarter-normalized link-weight normalizations

	Trans. only	Untreated only				Full sample			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Partner purchase growth $g^r_{j,t}$	−0.001	−0.001	−0.001	0.000	−0.000	−0.001	−0.001	0.000	−0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$T_j \times g^r_{j,t}$	0.001	0.001	0.001	0.001	0.002	0.001	0.001	0.000	0.001
	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.236***	0.231***	0.207***	0.220***	0.218***	0.256***	0.254***	0.257***	0.257***
	(0.047)	(0.045)	(0.051)	(0.056)	(0.055)	(0.038)	(0.038)	(0.038)	(0.038)
$T_i \times g^r_{j,t}$						0.002	0.001	0.001	0.000
						(0.003)	(0.002)	(0.003)	(0.002)
$T_i \times T_j \times g^r_{j,t}$						−0.003	−0.001	−0.003	−0.002
						(0.004)	(0.003)	(0.004)	(0.003)
Partner belief $\pi^{e,h}_{j,t}$		−0.011	−0.018	−0.031	−0.012	−0.011	−0.018	−0.034	−0.012
		(0.028)	(0.029)	(0.030)	(0.030)	(0.028)	(0.029)	(0.032)	(0.030)
$T_j \times \pi^{e,h}_{j,t}$		0.062**	0.056***	0.099***	0.062**	0.062**	0.055***	0.097***	0.058**
		(0.029)	(0.018)	(0.030)	(0.031)	(0.029)	(0.018)	(0.031)	(0.031)
$T_i \times \pi^{e,h}_{j,t}$						0.046	0.044**	0.075**	0.044
						(0.033)	(0.019)	(0.033)	(0.032)
$T_i \times T_j \times \pi^{e,h}_{j,t}$						−0.057	−0.050**	−0.061	−0.052
						(0.042)	(0.023)	(0.043)	(0.036)
Observations	108,187	108,187	38,452	38,452	38,452	108,187	108,187	108,187	108,187
R^2	0.073	0.074	0.069	0.059	0.059	0.073	0.073	0.074	0.073

Notes: Each column reports a dyad–quarter regression of firm i's h-quarter-ahead inflation expectation, $\pi^{e,h}_{i,t}$, on customer signals observed in the same quarter. For customers, $g^r_{j,t}$ is customer j's year-on-year purchase growth (winsorized at the 1st and 99th percentiles and expressed in percentage points). Partner beliefs are $\pi^{e,h}_{j,t}$. Column (1) includes only the outcome signal and its interaction with partner treatment T_j ("Trans. only"). Columns (2)–(5) add the partner belief signal (Untreated only) under four alternative within-firm link-weight normalizations: "brd" uses the baseline breadth/centrality weights within customer links, "unif" assigns uniform weights across observed customer links, "val" weights links by transaction value (taxable amount shares), and "shr" tilts transaction-value weights by the buyer's share of partner j's total sales in that quarter. Columns (1)–(5) restrict to untreated focal firms ($T_i = 0$). Columns (6)–(9) use the full sample and include T_i interactions (triple-interaction specifications). All specifications include the baseline controls and fixed effects used in the main text, use analytic weights, and report standard errors using the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

Table L.9: Linear size interactions ($h = 12$)

	Size = supplier-side share ($\widetilde{\ell}^{(j)}_{ij,t}$)		Size = buyer-side share ($\widetilde{\ell}^{(i)}_{ij,t}$)	
	Price only	Untreated only	Price only	Untreated only
$g^p_{j,t}$	-0.002	-0.002	-0.007*	-0.007*
	(0.005)	(0.005)	(0.003)	(0.003)
$T_j \times g^p_{j,t}$	0.008	0.008	0.015**	0.015**
	(0.007)	(0.007)	(0.006)	(0.006)
$\widetilde{\ell}_{ij,t}$	0.023*	0.023*	0.023	0.023
	(0.012)	(0.012)	(0.019)	(0.019)
$g^p_{j,t} \times \widetilde{\ell}_{ij,t}$	0.001	0.001	0.000	0.000
	(0.001)	(0.001)	(0.002)	(0.002)
$T_j \times \widetilde{\ell}_{ij,t}$	-0.028**	-0.028**	-0.025	-0.025
	(0.013)	(0.013)	(0.017)	(0.017)
$T_j \times g^p_{j,t} \times \widetilde{\ell}_{ij,t}$	-0.003*	-0.003*	-0.001	-0.001
	(0.001)	(0.001)	(0.002)	(0.002)
$\pi^{e,12}_{j,t}$	-0.025	-0.025	-0.014	-0.014
	(0.022)	(0.022)	(0.018)	(0.018)
$T_j \times \pi^{e,12}_{j,t}$	0.033*	0.033*	0.028**	0.028**
	(0.016)	(0.016)	(0.012)	(0.012)
$\pi^{e,12}_{j,t} \times \widetilde{\ell}_{ij,t}$	-0.011*	-0.011*	-0.003	-0.003
	(0.006)	(0.006)	(0.005)	(0.005)
$T_j \times \pi^{e,12}_{j,t} \times \widetilde{\ell}_{ij,t}$	0.016**	0.016**	-0.002	-0.002
	(0.007)	(0.007)	(0.004)	(0.004)
Lagged expectation $y_{ij,t-1}$	0.200**	0.200**	0.196***	0.196***
	(0.074)	(0.074)	(0.075)	(0.075)
Observations	44855		44855	
R^2	0.041		0.037	

Notes: Supplier relationships (firm j supplies firm i). Uniform link weights assign equal weight to each observed relationship within a buyer and quarter. Relationship size is the centered log transaction-value share computed prior to estimation. The dependent variable is the 12-quarters-ahead inflation expectation $\pi^{e,12}_{i,t}$. Sample is restricted to untreated buyers ($T_i = 0$). Standard errors follow the baseline procedure. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

by incentives to share information, whereas price-based learning operates more uniformly across relationships.

Implied marginal effects. Appendix Table L.10 translates the linear interaction coefficients into treated-supplier learning premia evaluated at P10 and P90 percentiles of the size distribution, as well as at the mean. The marginal-effects evidence confirms that belief-based learning at horizon $h = 12$ is strongly concentrated in economically important relationships. Under the supplier-side size definition, the treated-supplier belief premium is small or negligible at low-size relationships but rises sharply at the top of the size distribution. In contrast, the corresponding price-based premium displays much flatter size dependence and, if anything, weakly declines with size.

Table L.10: Treated-supplier learning premia at P10 and P90 ($h = 12$)

	P10	P90	$P90 - P10$	Mean
Panel A: Size = supplier-side share ($\widetilde{\ell}^{(j)}_{ij,t}$)				
Price signal $g^p_{j,t}$	0.019	-0.004	-0.023*	0.007
	(0.011)	(0.007)	(0.012)	(0.007)
Belief signal $\pi^{e,12}_{j,t}$	-0.034	0.107**	0.141**	0.038**
	(0.025)	(0.047)	(0.067)	(0.017)
Panel B: Size = buyer-side share ($\widetilde{\ell}^{(i)}_{ij,t}$)				
Price signal $g^p_{j,t}$	0.020*	0.011	-0.008	0.015**
	(0.011)	(0.010)	(0.016)	(0.006)
Belief signal $\pi^{e,12}_{j,t}$	0.018	0.036	0.018	0.027**
	(0.017)	(0.022)	(0.034)	(0.010)

Notes: Supplier relationships. Entries report the treated-supplier learning premium (incremental slope when $T_j = 1$) evaluated at the 10th (P10) and 90th (P90) percentiles of the relevant size distribution among untreated buyers, as well as the implied difference and the premium evaluated at the mean. Standard errors are delta-method standard errors computed from the estimated covariance matrix of the underlying $h = 12$ horse-race regression. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Relative to the pooled estimates, the single-horizon marginal effects are estimated with less precision, as expected given the smaller effective sample size. Nevertheless, the ranking of effects across size percentiles and across information channels remains intact.

Nonparametric bin regressions. Appendix Table L.11 provides a nonparametric check by estimating the horse-race specification separately within terciles of relationship size. The bin-level results are consistent with the parametric evidence. For supplier-side size, belief-based learning from treated suppliers is concentrated in the largest-size bin, while the effect is muted or imprecisely estimated in smaller relationships.

Price-based learning shows weaker and less monotonic variation across bins.

Customer-direction robustness. For completeness, Appendix Tables L.12, L.13, and L.14 report analogous results for the customer-direction sample. These estimates are less stable and economically less interpretable, reflecting the fact that the learning mechanism of interest operates primarily through suppliers. Importantly, none of the customer-direction results overturn

Table L.11: Treated-supplier premia by size bins ($h = 12$)

	Bin 1	Bin 2	Bin 3
Panel A: Size = supplier-side share bins ($\widetilde{\ell}^{(j)}_{ij,t}$)			
$T_j \times g^p_{j,t}$	-0.003	-0.007*	0.066**
	(0.016)	(0.004)	(0.031)
$T_j \times \pi^{e,12}_{j,t}$	0.052	0.015	0.042
	(0.060)	(0.015)	(0.098)
Observations	14,604	14,763	15,367
R^2	0.056	0.056	0.059
Panel B: Size = buyer-side share bins ($\widetilde{\ell}^{(i)}_{ij,t}$)			
$T_j \times g^p_{j,t}$	0.019	0.007	0.007
	(0.021)	(0.009)	(0.015)
$T_j \times \pi^{e,12}_{j,t}$	0.068	-0.003	0.009
	(0.050)	(0.021)	(0.063)
Observations	15,137	15,041	14,565
R^2	0.032	0.051	0.042

Notes: Supplier relationships (firm j supplies firm i). Each column reports the treated-supplier premium from the $h = 12$ horse-race regression estimated within the corresponding tercile of relationship size (bins defined using untreated buyers and the relevant pre-estimation size measure). Coefficients shown are the incremental slopes on supplier price growth ($T_j \times g^p_{j,t}$) and on supplier beliefs ($T_j \times \pi^{e,12}_{j,t}$). Uniform within-buyer–quarter weights are used. Standard errors in parentheses follow the baseline procedure. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

the core conclusion that belief-based learning heterogeneity is strongest in relationships where incentives to share information are high.

L.2 Robustness: $h = 12$ evidence on supplier breadth

This appendix examines the robustness of the breadth-heterogeneity results in Section 4 when attention is restricted to a single forecast horizon. As in Appendix L.1, we re-estimate the uniform-weight specifications using only the 12-quarters-ahead inflation expectation, $\pi^{e,12}_{i,t}$, rather than pooling across horizons.

Design. The empirical design mirrors the pooled analysis as closely as possible. The dependent variable is firm i's 12-quarter-ahead inflation expectation reported in quarter t, denoted $y_{ij,t} \equiv \pi^{e,12}_{i,t}$. There is no horizon stacking and hence no horizon fixed effects. The lag structure remains horizon-specific and we include the appropriate lagged outcome $y_{ij,t-1} = \pi^{e,12}_{i,t-1}$. All specifications focus on supplier relationships (links in which firm j supplies firm i), restrict attention to untreated buyers ($T_i = 0$), use uniform within-buyer–quarter weights, and include the same fixed effects and controls as in the pooled benchmark.

Breadth is measured using SIGE upstream degree. *Partner breadth* is defined as the number of upstream suppliers of firm j observed in quarter t, while *own breadth* maps the same object to firm i. In both cases, we work with centered log breadth constructed prior to estimation.

Table L.12: Linear size interactions ($h = 12$)

	Size = supplier-side share ($\widetilde{\ell}^{(j)}_{ij,t}$)		Size = buyer-side share ($\widetilde{\ell}^{(i)}_{ij,t}$)	
	Price only	Untreated only	Price only	Untreated only
$g^p_{j,t}$	0.004	0.004	0.005	0.005
	(0.005)	(0.005)	(0.006)	(0.006)
$T_j \times g^p_{j,t}$	-0.004	-0.004	-0.002	-0.002
	(0.008)	(0.008)	(0.007)	(0.007)
$\widetilde{\ell}_{ij,t}$	0.003	0.003	0.004	0.004
	(0.022)	(0.022)	(0.017)	(0.017)
$g^p_{j,t} \times \widetilde{\ell}_{ij,t}$	0.002	0.002	-0.000	-0.000
	(0.002)	(0.002)	(0.002)	(0.002)
$T_j \times \widetilde{\ell}_{ij,t}$	0.005	0.005	-0.002	-0.002
	(0.031)	(0.031)	(0.016)	(0.016)
$T_j \times g^p_{j,t} \times \widetilde{\ell}_{ij,t}$	-0.001	-0.001	0.000	0.000
	(0.002)	(0.002)	(0.002)	(0.002)
$\pi^{e,12}_{j,t}$	-0.018	-0.018	-0.026*	-0.026*
	(0.022)	(0.022)	(0.015)	(0.015)
$T_j \times \pi^{e,12}_{j,t}$	0.090***	0.090***	0.069***	0.069***
	(0.027)	(0.027)	(0.024)	(0.024)
$\pi^{e,12}_{j,t} \times \widetilde{\ell}_{ij,t}$	-0.020	-0.020	-0.009	-0.009
	(0.014)	(0.014)	(0.006)	(0.006)
$T_j \times \pi^{e,12}_{j,t} \times \widetilde{\ell}_{ij,t}$	0.007	0.007	0.007	0.007
	(0.012)	(0.012)	(0.009)	(0.009)
Lagged expectation $y_{ij,t-1}$	0.197***	0.197***	0.198***	0.198***
	(0.075)	(0.075)	(0.074)	(0.074)
Observations	44855		44855	
R^2	0.041		0.041	

Notes: Customer relationships (firm i supplies firm j; i.e., links where j is a customer of i). Uniform link weights assign equal weight to each observed relationship within a buyer and quarter. Relationship size is the centered log transaction-value share computed prior to estimation. The dependent variable is the 12-quarters-ahead inflation expectation $\pi^{e,12}_{i,t}$. Sample is restricted to untreated buyers ($T_i = 0$). Standard errors follow the baseline procedure. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table L.13: Treated-customer learning premia at P10 and P90 ($h = 12$)

	P10	P90	$P90 - P10$	Mean
Panel A: Size = supplier-side share ($\widetilde{\ell}^{(j)}_{ij,t}$)				
Price signal $g^{p}_{j,t}$	0.007	-0.014	-0.021	-0.004
	(0.019)	(0.012)	(0.028)	(0.007)
Belief signal $\pi^{e,12}_{j,t}$	0.061*	0.122**	0.061	0.093***
	(0.032)	(0.056)	(0.074)	(0.028)
Panel B: Size = buyer-side share ($\widetilde{\ell}^{(i)}_{ij,t}$)				
Price signal $g^{p}_{j,t}$	0.002	-0.006	-0.008	-0.002
	(0.016)	(0.009)	(0.021)	(0.008)
Belief signal $\pi^{e,12}_{j,t}$	0.068***	0.089***	0.021	0.078***
	(0.025)	(0.030)	(0.043)	(0.022)

Notes: Customer relationships. Entries report the treated-partner learning premium (incremental slope when $T_j = 1$) evaluated at the 10th (P10) and 90th (P90) percentiles of the relevant size distribution among untreated buyers, as well as the implied difference and the premium evaluated at the mean. Standard errors are delta-method standard errors computed from the estimated covariance matrix of the underlying $h = 12$ horse-race regression. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table L.14: Treated-partner premia by size bins ($h = 12$)

	Bin 1	Bin 2	Bin 3
Panel A: Size = supplier-side share bins ($\widetilde{\ell}^{(j)}_{ij,t}$)			
$T_j \times g^{p}_{j,t}$	-0.029	-0.024*	0.004
	(0.023)	(0.012)	(0.071)
$T_j \times \pi^{e,12}_{j,t}$	0.150***	0.089*	-0.208
	(0.039)	(0.048)	(0.139)
Observations	16,700	16,189	16,016
R^2	0.087	0.021	0.072
Panel B: Size = buyer-side share bins ($\widetilde{\ell}^{(i)}_{ij,t}$)			
$T_j \times g^{p}_{j,t}$	0.002	0.002	0.007
	(0.012)	(0.010)	(0.042)
$T_j \times \pi^{e,12}_{j,t}$	0.058	-0.008	0.094**
	(0.039)	(0.015)	(0.045)
Observations	16,847	16,351	15,743
R^2	0.049	0.080	0.044

Notes: Customer-direction robustness sample (as defined in the corresponding regressions). Each column reports the treated-partner premium from the $h = 12$ horse-race regression estimated within the corresponding tercile of relationship size. Coefficients shown are the incremental slopes on partner price growth ($T_j \times g^{p}_{j,t}$) and on partner beliefs ($T_j \times \pi^{e,12}_{j,t}$). Uniform within-firm–quarter weights are used. Standard errors in parentheses follow the baseline procedure. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Linear interaction estimates. Appendix Table L.15 reports linear interaction estimates for the $h = 12$ outcome, separately for partner breadth and own breadth, and for price-only versus horse-race specifications. The qualitative patterns closely mirror those in the pooled analysis. As partner breadth increases, the treated-supplier premium in the belief channel declines, while the corresponding price-channel premium increases. This configuration is consistent with breadth capturing network complexity: suppliers with many upstream partners aggregate more information into prices, but face higher costs of transmitting detailed belief information to downstream buyers.

Table L.15: Linear breadth interactions ($h = 12$)

	Partner breadth ($\widetilde{b}^{(j)}_{j,t}$)		Own breadth ($\widetilde{b}^{(i)}_{i,t}$)	
	Price only	Untreated only	Price only	Untreated only
$g^p_{j,t}$	0.002	0.002	0.002	0.002
	(0.007)	(0.007)	(0.004)	(0.004)
$T_j \times g^p_{j,t}$	0.004	0.004	0.001	0.001
	(0.009)	(0.009)	(0.003)	(0.003)
$\widetilde{b}_{\cdot,t}$	-0.089**	-0.089**	0.028	0.028
	(0.041)	(0.041)	(0.153)	(0.153)
$g^p_{j,t} \times \widetilde{b}_{\cdot,t}$	-0.011*	-0.011*	0.009	0.009
	(0.005)	(0.005)	(0.006)	(0.006)
$T_j \times \widetilde{b}_{\cdot,t}$	0.088*	0.088*	0.064	0.064
	(0.049)	(0.049)	(0.062)	(0.062)
$T_j \times g^p_{j,t} \times \widetilde{b}_{\cdot,t}$	0.010**	0.010**	-0.013	-0.013
	(0.004)	(0.004)	(0.009)	(0.009)
$\pi^{e,12}_{j,t}$	-0.032	-0.032	0.006	0.006
	(0.027)	(0.027)	(0.064)	(0.064)
$T_j \times \pi^{e,12}_{j,t}$	0.043*	0.043*	0.006	0.006
	(0.022)	(0.022)	(0.021)	(0.021)
$\pi^{e,12}_{j,t} \times \widetilde{b}_{\cdot,t}$	0.033**	0.033**	-0.003	-0.003
	(0.012)	(0.012)	(0.010)	(0.010)
$T_j \times \pi^{e,12}_{j,t} \times \widetilde{b}_{\cdot,t}$	-0.032**	-0.032**	0.002	0.002
	(0.012)	(0.012)	(0.003)	(0.003)
Lagged expectation $y_{ij,t-1}$	0.199**	0.199**	0.203***	0.203***
	(0.074)	(0.074)	(0.073)	(0.073)
Observations	44855		44855	
R^2	0.041		0.042	

Notes: Supplier relationships (firm j supplies firm i). Breadth is log(Breadth), where Breadth is the SIGE upstream degree (*j_n_supplall*), demeaned within the estimation sample. "Partner breadth" uses the upstream degree of partner j in quarter t; "own breadth" maps the same object to firm i by (i, t). The dependent variable is the 12-quarters-ahead inflation expectation $\pi^{e,12}_{i,t}$ and the sample is restricted to untreated firms ($T_i = 0$). Uniform link weights assign equal weight to each observed link within a firm and quarter. Standard errors follow the baseline procedure. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

By contrast, heterogeneity with respect to own breadth is weaker and less stable, suggesting that it is the information environment of the supplier, rather than that of the buyer, that primarily governs the composition of learning.

Implied marginal effects. Appendix Table L.16 translates the linear interaction coefficients into treated-supplier learning premia evaluated at P10 and P90 percentiles of the breadth distribution, as well as at the mean. The marginal-effects evidence confirms that breadth heterogeneity operates in the opposite direction from size heterogeneity. For partner breadth, the treated-supplier belief premium is large and positive at low breadth but declines sharply and becomes small or negative at high breadth. In contrast, the treated-supplier price premium increases with breadth, indicating that prices become the dominant information carrier in complex upstream networks.

Relative to the pooled estimates, the single-horizon marginal effects are estimated with less precision, but the ranking of effects across percentiles and across information channels remains intact.

Table L.16: Treated-supplier learning premia at P10 and P90 by breadth ($h = 12$)

	P10	P90	$P90 - P10$	Mean
Panel A: Partner breadth ($\widetilde{b}^{(j)}_{j,t}$)				
Price signal $g^{p}_{j,t}$	-0.009	0.016^{*}	0.024^{**}	0.003
	(0.007)	(0.010)	(0.009)	(0.007)
Belief signal $\pi^{e,12}_{j,t}$	0.074^{***}	-0.009	-0.083^{**}	0.036^{*}
	(0.032)	(0.018)	(0.031)	(0.021)
Panel B: Own breadth ($\widetilde{b}^{(i)}_{i,t}$)				
Price signal $g^{p}_{j,t}$	-0.006	0.012	0.018	0.002
	(0.008)	(0.011)	(0.014)	(0.007)
Belief signal $\pi^{e,12}_{j,t}$	0.031	0.019	-0.012	0.025^{*}
	(0.021)	(0.020)	(0.028)	(0.013)

Notes: Supplier relationships. Entries report the treated-supplier learning premium (incremental slope when $T_j = 1$) evaluated at the 10th (P10) and 90th (P90) percentiles of the relevant breadth distribution among untreated buyers, as well as the implied difference and the premium evaluated at the mean. Standard errors are delta-method standard errors computed from the estimated covariance matrix of the underlying $h = 12$ horse-race regression. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Nonparametric bin regressions. Appendix Table L.17 provides a nonparametric check by estimating the horse-race specification separately within terciles of breadth. The bin-level evidence reinforces the parametric results. For partner breadth, belief-based learning from treated suppliers is strongest in the lowest-breadth bin and declines monotonically across bins, while price-based learning becomes stronger at higher breadth. Results based on own breadth are weaker and less monotonic.

Customer-direction robustness. Appendix Tables L.18, L.19, and L.20 report analogous results for the customer-direction sample. As in the pooled analysis, these estimates are less stable and economically less interpretable. Importantly, none of the customer-direction results overturn the central conclusion that greater supplier breadth shifts learning away from beliefs and toward prices.

Symmetric price signal. Table L.21 reports the pooled-horizon joint-direction specification using survey-reported own price growth $g^{p}_{j,t}$ for *both* upstream and downstream links, replacing

Table L.17: Treated-supplier premia by breadth bins ($h = 12$)

	Bin 1	Bin 2	Bin 3
Panel A: Partner breadth bins			
$T_j \times g^p_{j,t}$	−0.009	0.008	0.017
	(0.011)	(0.006)	(0.013)
$T_j \times \pi^{e,12}_{j,t}$	0.077***	0.026	−0.002
	(0.025)	(0.018)	(0.015)
Panel B: Own breadth bins			
$T_j \times g^p_{j,t}$	−0.004	0.003	0.006
	(0.010)	(0.008)	(0.012)
$T_j \times \pi^{e,12}_{j,t}$	0.036	0.019	0.011
	(0.028)	(0.020)	(0.024)

Notes: Supplier relationships (firm j supplies firm i). Each column reports the treated-supplier premium from the $h = 12$ horse-race regression estimated within the corresponding tercile of breadth (bins defined using untreated buyers and the relevant pre-estimation breadth measure). Coefficients shown are the incremental slopes on supplier price growth ($T_j \times g^p_{j,t}$) and on supplier beliefs ($T_j \times \pi^{e,12}_{j,t}$). Uniform within-buyer–quarter weights are used. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

the administrative purchase growth $g^r_{j,t}$ used for customers in the baseline. This specification tests whether the downstream learning results are driven by the particular choice of customer outcome signal.

Table L.18: Linear breadth interactions ($h = 12$)

	Partner breadth ($\widetilde{b}^{(j)}_{j,t}$)		Own breadth ($\widetilde{b}^{(i)}_{i,t}$)	
	Price only	Untreated only	Price only	Untreated only
$g^p_{j,t}$	0.006	0.006	0.005	0.005
	(0.005)	(0.005)	(0.007)	(0.007)
$T_j \times g^p_{j,t}$	-0.005	-0.005	-0.001	-0.001
	(0.006)	(0.006)	(0.009)	(0.009)
$\widetilde{b}_{\cdot,t}$	-0.070	-0.070	0.048	0.048
	(0.074)	(0.074)	(0.221)	(0.221)
$g^p_{j,t} \times \widetilde{b}_{\cdot,t}$	-0.005	-0.005	0.003	0.003
	(0.005)	(0.005)	(0.006)	(0.006)
$T_j \times \widetilde{b}_{\cdot,t}$	0.068	0.068	-0.098*	-0.098*
	(0.092)	(0.092)	(0.054)	(0.054)
$T_j \times g^p_{j,t} \times \widetilde{b}_{\cdot,t}$	0.009	0.009	0.003	0.003
	(0.009)	(0.009)	(0.006)	(0.006)
$\pi^{e,12}_{j,t}$	-0.064**	-0.064**	-0.056***	-0.056***
	(0.031)	(0.031)	(0.017)	(0.017)
$T_j \times \pi^{e,12}_{j,t}$	0.055**	0.055**	0.057***	0.057***
	(0.022)	(0.022)	(0.017)	(0.017)
$\pi^{e,12}_{j,t} \times \widetilde{b}_{\cdot,t}$	0.051**	0.051**	0.022**	0.022**
	(0.021)	(0.021)	(0.009)	(0.009)
$T_j \times \pi^{e,12}_{j,t} \times \widetilde{b}_{\cdot,t}$	-0.051**	-0.051**	-0.024***	-0.024***
	(0.021)	(0.021)	(0.009)	(0.009)
Lagged expectation $y_{ij,t-1}$	0.193***	0.193***	0.193***	0.193***
	(0.051)	(0.051)	(0.051)	(0.051)
Observations	49050		49050	
R^2	0.050		0.050	

Notes: Customer relationships (firm i supplies firm j). Breadth is log(Breadth), where Breadth is the SIGE upstream degree (*j_n_supplall*), demeaned within the estimation sample. "Partner breadth" uses the upstream degree of partner j in quarter t; "own breadth" maps the same object to firm i by (i, t). The dependent variable is the 12-quarters-ahead inflation expectation $\pi^{e,12}_{i,t}$ and the sample is restricted to untreated firms ($T_i = 0$). Uniform link weights assign equal weight to each observed link within a firm and quarter. Standard errors follow the baseline procedure. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table L.19: Treated-partner learning premia at P10 and P90 by breadth ($h = 12$)

	P10	P90	$P90 - P10$	Mean
Panel A: Partner breadth ($\widetilde{b}^{(j)}_{j,t}$)				
Price signal $g^p_{j,t}$	−0.016	0.006	0.021	−0.006
	(0.010)	(0.014)	(0.021)	(0.006)
Belief signal $\pi^{e,12}_{j,t}$	0.139**	0.052*	−0.087	0.098***
	(0.058)	(0.026)	(0.070)	(0.030)
Panel B: Own breadth ($\widetilde{b}^{(i)}_{i,t}$)				
Price signal $g^p_{j,t}$	−0.004	0.003	0.007	−0.000
	(0.008)	(0.015)	(0.015)	(0.010)
Belief signal $\pi^{e,12}_{j,t}$	0.098**	0.113***	0.014	0.106***
	(0.038)	(0.035)	(0.061)	(0.020)

Notes: Customer relationships (firm i supplies firm j). Entries report the treated-partner learning premium (incremental slope when $T_j = 1$) evaluated at the 10th (P10) and 90th (P90) percentiles of the relevant breadth distribution among untreated firms, as well as the implied difference and the premium evaluated at the mean. Standard errors are delta-method standard errors computed from the estimated covariance matrix of the underlying $h = 12$ horse-race regression. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table L.20: Treated-partner premia by breadth bins ($h = 12$)

	Bin 1	Bin 2	Bin 3
Panel A: Partner breadth bins			
$T_j \times g^p_{j,t}$	0.021	0.006	-0.029
	(0.053)	(0.010)	(0.031)
$T_j \times \pi^{e,12}_{j,t}$	0.070	0.033	0.012
	(0.059)	(0.034)	(0.110)
Observations	16,402	16,361	16,168
R^2	0.069	0.042	0.028
Panel B: Own breadth bins			
$T_j \times g^p_{j,t}$	0.013	-0.009	-0.001
	(0.017)	(0.008)	(0.001)
$T_j \times \pi^{e,12}_{j,t}$	0.174***	-0.014	-0.011
	(0.032)	(0.023)	(0.008)
Observations	15,163	16,915	16,969
R^2	0.048	0.061	0.886

Notes: Customer relationships (firm i supplies firm j). Each column reports the treated-partner premium from the $h = 12$ horse-race regression estimated within the corresponding tercile of breadth (bins defined using untreated firms and the relevant pre-estimation breadth measure). Coefficients shown are the incremental slopes on partner price growth ($T_j \times g^p_{j,t}$) and on partner beliefs ($T_j \times \pi^{e,12}_{j,t}$). Uniform within-firm–quarter weights are used. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

Table L.21: Learning from supply-chain partners: symmetric price signal

	Untreated only ($T_{i,t}=0$)				Full sample			
	(1) unif	(2) brd	(3) val	(4) shr	(5) unif	(6) brd	(7) val	(8) shr
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^p_{j,t}$)								
β_0: $g_{j,t}$	−0.002	−0.003	−0.002	0.007	−0.004	−0.006	−0.006	0.006
	(0.005)	(0.004)	(0.004)	(0.009)	(0.005)	(0.004)	(0.005)	(0.009)
β_1: $T_{j,t}\, g_{j,t}$	0.009	0.009*	0.000	−0.003	0.012	0.013**	0.004	−0.002
	(0.008)	(0.005)	(0.005)	(0.007)	(0.008)	(0.006)	(0.006)	(0.007)
δ_0: $\pi^{e,h}_{j,t}$	0.002	0.013	−0.003	0.008	−0.008	0.010	−0.017	0.016
	(0.019)	(0.018)	(0.031)	(0.029)	(0.026)	(0.018)	(0.029)	(0.027)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.022	0.028	0.056*	0.037	0.034**	0.022	0.065**	0.031*
	(0.013)	(0.020)	(0.030)	(0.022)	(0.016)	(0.016)	(0.028)	(0.017)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					−0.035*	−0.027	−0.042*	−0.003
					(0.017)	(0.020)	(0.022)	(0.013)
Panel B: Downstream signals ($g_{j,t}$ = customer price growth $g^p_{j,t}$)								
β_0: $g_{j,t}$	−0.002	0.001	−0.003	−0.017*	−0.003	−0.001	−0.003	−0.020*
	(0.006)	(0.003)	(0.005)	(0.010)	(0.007)	(0.005)	(0.006)	(0.011)
β_1: $T_{j,t}\, g_{j,t}$	0.003	−0.002	0.003	0.019**	0.002	−0.001	0.001	0.019*
	(0.005)	(0.006)	(0.006)	(0.009)	(0.006)	(0.006)	(0.006)	(0.010)
δ_0: $\pi^{e,h}_{j,t}$	−0.011	−0.008	−0.017	0.013	0.004	−0.007	0.002	0.015
	(0.022)	(0.023)	(0.032)	(0.025)	(0.019)	(0.028)	(0.029)	(0.023)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.065***	0.070***	0.091**	0.045*	0.040**	0.076**	0.070**	0.034
	(0.020)	(0.024)	(0.034)	(0.025)	(0.016)	(0.034)	(0.033)	(0.029)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					−0.037	−0.061	−0.041	−0.029
					(0.026)	(0.050)	(0.044)	(0.037)
ρ: $\pi^{e,h}_{i,t-1}$	0.243***	0.255***	0.255***	0.255***	0.273***	0.280***	0.279***	0.279***
	(0.062)	(0.055)	(0.056)	(0.056)	(0.050)	(0.051)	(0.051)	(0.051)
Observations	187,926	187,926	187,926	187,926	558,932	558,932	558,932	558,932
R^2	0.061	0.068	0.068	0.068	0.075	0.080	0.081	0.080

Notes: Same specification as Table 3, except that the downstream outcome signal is customer j's survey-reported own price growth $g^p_{j,t}$ instead of administrative purchase growth $g^r_{j,t}$. Both panels therefore use the same signal type. The sample is slightly larger because price growth ($g^p_{j,t}$) has fewer missing values than administrative purchase growth. For brevity, β_2, β_3 (outcome channel, own treatment), and δ_2 (belief channel, own treatment) are estimated but suppressed. All other details as in Table 3. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

The belief-channel results are qualitatively unchanged. Downstream, δ_1 ranges from 0.045 to 0.091 across untreated-only specifications (Panel B, columns 1–4), compared with 0.064 to 0.113 under the baseline purchase-growth signal (Table 3). Upstream coefficients are similarly close to baseline. The outcome channel downstream is weak and imprecise under most weighting schemes, with the exception of tilted weights (column 4), where $\beta_1 = 0.019$ ($p < 0.05$) suggests that customer price growth carries some information in high-intensity relationships.

Robustness to partner-level controls. Table L.22 replicates the main specification of Table 3 adding partner-level controls: j's sector×quarter, size class×quarter, and geographic area×quarter fixed effects. These absorb time-varying differences in partner observables that might correlate with treatment status. All belief-channel coefficients are virtually unchanged, confirming that the treatment contrast is not driven by systematic differences in partner characteristics.

Table L.22: Robustness: partner-level controls

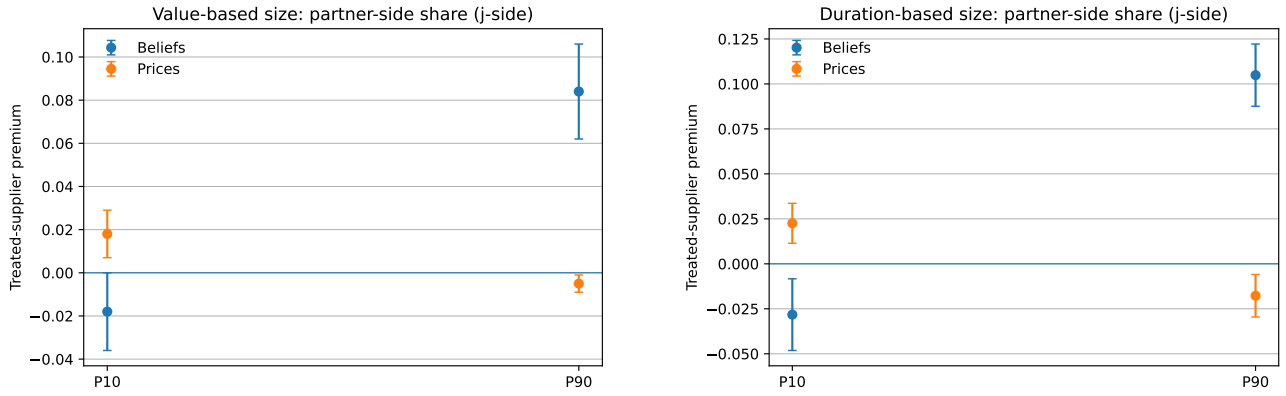
	Untreated only ($T_{i,t} = 0$)				Full sample			
	(1) brd	(2) unif	(3) val	(4) shr	(5) brd	(6) unif	(7) val	(8) shr
Panel A: Upstream signals ($g_{j,t} = $ supplier price growth $g^p_{j,t}$)								
β_0: $g^p_{j,t}$	−0.010*	−0.005	−0.005	0.003	−0.013*	−0.009	−0.009	0.004
	(0.005)	(0.005)	(0.005)	(0.009)	(0.006)	(0.006)	(0.007)	(0.010)
β_1: $T_{j,t}\, g^p_{j,t}$	0.017**	0.012	0.004	0.001	0.021**	0.017*	0.008	0.001
	(0.007)	(0.009)	(0.005)	(0.007)	(0.008)	(0.009)	(0.007)	(0.007)
δ_0: $\pi^{e,h}_{j,t}$	0.015	0.001	−0.015	0.004	0.022	−0.001	−0.008	0.018
	(0.016)	(0.021)	(0.031)	(0.026)	(0.020)	(0.026)	(0.031)	(0.029)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.031	0.025	0.064*	0.034*	0.020	0.032*	0.061**	0.030
	(0.020)	(0.018)	(0.031)	(0.018)	(0.016)	(0.017)	(0.028)	(0.019)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					−0.030	−0.032	−0.043*	−0.002
					(0.020)	(0.020)	(0.024)	(0.019)
Panel B: Downstream signals ($g_{j,t} = $ customer purchase growth $g^r_{j,t}$)								
β_0: $g^r_{j,t}$	−0.001	−0.001	−0.000	−0.001	0.001	0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)
β_1: $T_{j,t}\, g^r_{j,t}$	0.001	0.002	0.000	0.003***	−0.001	0.001	−0.002	0.001
	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
δ_0: $\pi^{e,h}_{j,t}$	−0.026	−0.033	−0.046	−0.000	−0.025	−0.017	−0.026	−0.003
	(0.026)	(0.030)	(0.034)	(0.029)	(0.018)	(0.013)	(0.019)	(0.022)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.096***	0.081***	0.119***	0.055*	0.091***	0.057***	0.085***	0.045
	(0.026)	(0.025)	(0.033)	(0.026)	(0.029)	(0.017)	(0.029)	(0.031)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					−0.098**	−0.067**	−0.066	−0.048
					(0.044)	(0.025)	(0.042)	(0.038)
ρ: $\pi^{e,h}_{i,t-1}$	0.265***	0.252***	0.260***	0.261***	0.289***	0.280***	0.289***	0.290***
	(0.050)	(0.057)	(0.051)	(0.052)	(0.045)	(0.046)	(0.045)	(0.045)
Observations	166,370	166,370	166,370	166,370	497,406	497,406	497,406	497,406
R^2	0.071	0.063	0.069	0.069	0.083	0.078	0.083	0.083

Notes: Same specification as Table 3, with the addition of partner-level fixed effects: j's sector×quarter, size class×quarter, and geographic area×quarter. For brevity, β_2, β_3 (outcome channel, own treatment), and δ_2 (belief channel, own treatment) are estimated but suppressed. All other details as in Table 3. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

M Own-side heterogeneity

This appendix reports the full own-side heterogeneity regression estimates that underlie the own-side marginals in the main text (Table M.1). All specifications use equation (7) with uniform link weights, pooled horizons $h \in \{6, 12\}$, and the same fixed effects and controls as in the main text.

Figure M.1: Treated-partner learning premia at P10 and P90: partner-side value vs persistence



Notes: The figure plots treated-partner learning premia (incremental slopes when $T_{j,t} = 1$) evaluated at P10 and P90 of the partner-side size distribution among untreated buyers. The left panel uses value-based shares (transaction values cumulated over the sample); the right panel uses persistence-based shares (relationship persistence). For each panel, we report premia for the upstream outcome channel (supplier price growth) and the upstream belief channel (supplier inflation expectations). Error bars represent one standard error (delta-method standard errors for the P10/P90 evaluations).

Table M.2 reports the full coefficient estimates. The key contrast with the partner-side results (Table 4) is the upstream belief channel: the treated-partner size interaction δ_3 is negative (-0.003, s.e. 0.003) under own-side value-based size, compared with positive (0.011, s.e. 0.003) under partner-side size. The persistence-based interaction is insignificant under own-side shares, in contrast to the large and significant partner-side gradient (0.035, s.e. 0.006). Own breadth shows a positive but imprecise upstream belief interaction (0.024, s.e. 0.017), while own-breadth downstream interactions are negative and marginally significant (-0.015, s.e. 0.008), suggesting that firms with more downstream connections learn less from each individual customer's beliefs.

When value and persistence are entered jointly (column 4), neither margin produces a significant gradient for upstream beliefs. Table M.3 reports Wald tests: neither value nor persistence margins are individually significant ($p_{\text{val}} = 0.75$, $p_{\text{per}} = 0.49$), and the joint test is not significant ($p = 0.58$). For the downstream belief channel, the persistence margin is marginally significant ($p_{\text{per}} = 0.09$) but neither the value margin ($p_{\text{val}} = 0.41$) nor the joint test ($p = 0.22$) reaches the 5% level.

Table M.1: Treated-partner learning premia at P10 and P90 (own-side)

	P10	P90	Δ(P90–P10)
Panel A: Value-based size $\widetilde{\ell}_{ij}^{(i),\mathrm{val}}$			
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)	0.040*	0.017	−0.023
	(0.020)	(0.016)	(0.025)
$T_{j,t} \times g_{j,t}^{p}$	0.028	0.008	−0.020
	(0.020)	(0.010)	(0.023)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.037	0.090***	0.054
	(0.022)	(0.028)	(0.038)
$T_{j,t} \times g_{j,t}^{r}$	0.003	0.001	−0.001
	(0.003)	(0.002)	(0.003)
Panel B: Persistence $\widetilde{\ell}_{ij}^{(i),\mathrm{per}}$			
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)	0.025	0.022	−0.002
	(0.027)	(0.015)	(0.031)
$T_{j,t} \times g_{j,t}^{p}$	0.024	0.011	−0.013
	(0.022)	(0.009)	(0.025)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.031	0.087***	0.056
	(0.034)	(0.026)	(0.045)
$T_{j,t} \times g_{j,t}^{r}$	0.003	0.001	−0.002
	(0.003)	(0.003)	(0.006)
Panel C: Own breadth $\widetilde{b}_{i,t}^{(i)}$			
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)	0.021	0.134*	0.113
	(0.016)	(0.067)	(0.077)
$T_{j,t} \times g_{j,t}^{p}$	0.011	−0.050	−0.061
	(0.009)	(0.039)	(0.047)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.085***	0.015	−0.070*
	(0.024)	(0.020)	(0.036)
$T_{j,t} \times g_{j,t}^{r}$	0.001	0.005	0.004
	(0.002)	(0.003)	(0.003)

Notes: Entries evaluate the treated-partner premium at the 10th and 90th percentiles of the centered log own-side heterogeneity variable. Panels A, B, and C correspond to the own-side counterparts of the partner-side results in Table 5. Full own-side regression estimates are reported in Table M.2. Standard errors by the delta method. *, **, *** denote significance at the 10%, 5%, and 1% levels.

Table M.2: Heterogeneity in learning from partners: own-side

	(1) Value $\widetilde{\ell}_{ij}^{(i),\text{val}}$	(2) Persistence $\widetilde{\ell}_{ij}^{(i),\text{per}}$	(3) Breadth $\widetilde{b}_{i,t}^{(i)}$	(4) Val. + Pers. Joint
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^p_{j,t}$)				
β_1: $T_{j,t}\, g_{j,t}$	0.017	0.017	−0.014	0.020
	(0.010)	(0.011)	(0.011)	(0.015)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{val}}$	−0.002			−0.002
	(0.003)			(0.002)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{per}}$		−0.003		−0.003
		(0.006)		(0.007)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{b}_{i,t}^{(i)}$			−0.013	
			(0.010)	
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.028**	0.023	0.068***	0.035**
	(0.013)	(0.014)	(0.023)	(0.014)
δ_3: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{val}}$	−0.003			−0.001
	(0.003)			(0.003)
δ_3: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{per}}$		−0.001		−0.007
		(0.008)		(0.010)
δ_3: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{b}_{i,t}^{(i)}$			0.024	
			(0.017)	
Panel B: Downstream signals ($g_{j,t}$ = customer purchase growth $g^r_{j,t}$)				
β_1: $T_{j,t}\, g_{j,t}$	0.002	0.002**	0.003	0.002*
	(0.002)	(0.001)	(0.002)	(0.001)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{val}}$	−0.000			−0.000
	(0.000)			(0.001)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{per}}$		−0.001		0.000
		(0.001)		(0.002)
β_3: $T_{j,t}\, g_{j,t} \times \widetilde{b}_{i,t}^{(i)}$			0.001	
			(0.001)	
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.065***	0.060***	0.056***	0.053***
	(0.017)	(0.020)	(0.014)	(0.017)
δ_3: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{val}}$	0.006			0.002
	(0.004)			(0.003)
δ_3: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{\ell}_{ij}^{(i),\text{per}}$		0.015		0.017*
		(0.012)		(0.009)
δ_3: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{b}_{i,t}^{(i)}$			−0.015*	
			(0.008)	
ρ: $\pi^{e,h}_{i,t-1}$	0.248***	0.248***	0.248***	0.248***
	(0.061)	(0.061)	(0.061)	(0.061)
Observations	166,370	166,370	166,370	166,370
R^2	0.066	0.067	0.067	0.068

Notes: Each column estimates equation (7) on untreated firms ($T_{i,t} = 0$) with uniform link weights and pooled horizons $h \in \{6, 12\}$. The heterogeneity variable is: column (1) the centered log of partner j's share in firm i's total transaction value cumulated over the sample; column (2) the centered log of the dyad's persistence share in firm i's total relationship persistence; column (3) the centered log of firm i's direction-specific number of trading partners; column (4) value and persistence entered jointly. Since both panels are extracted from a single pooled-direction regression, ρ and fit statistics are common across directions. For brevity, only $T_{j,t}$-interaction terms and their heterogeneity gradients are shown. Marginal effects at P10 and P90 are reported in the main text (Table M.1).

Table M.3: Wald tests for joint and individual size-interaction significance (own-side)

	p-values			
Signal	Joint	Value only	Persistence only	Equality
$\pi^{e,h}_{j,t}$ (upstream)	0.583	0.752	0.489	0.613
$\pi^{e,h}_{j,t}$ (downstream)	0.219	0.410	0.091	0.129
$g^{p}_{j,t}$	0.583	0.307	0.715	0.898
$g^{r}_{j,t}$	0.889	0.775	0.973	0.922

Notes: Wald tests from the joint specification (column 4) of Table M.2.

N Value-based size: partner-side full specification

This appendix reports the full partner-side value-based size specification referenced in Section 4. Table 4 column (1) reports the $T_{j,t}$-interaction terms; the table below shows all estimated coefficients. Own-side value-based size results are reported in Table M.2 column (1).

Table N.1: Pooled-direction value-based size interactions: partner-side share $\widetilde{\ell}_{ij}^{(j),\mathrm{val}}$

	(1) Outcome only	(2) Outcome + beliefs
Common heterogeneity controls		
$\widetilde{\ell}_{ij}^{(j),\mathrm{val}}$	−0.014*	0.001
	(0.007)	(0.006)
$T_{j,t} \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$	0.025*	−0.001
	(0.013)	(0.008)
Upstream (supplier) signals		
$g_{j,t}^{p}$	−0.000	−0.001
	(0.004)	(0.004)
$T_{j,t} \times g_{j,t}^{p}$	0.005	0.006
	(0.006)	(0.006)
$g_{j,t}^{p} \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$	−0.000	0.000
	(0.001)	(0.001)
$T_{j,t} \times g_{j,t}^{p} \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$	−0.001	−0.003*
	(0.001)	(0.001)
$\pi_{j,t}^{e,h}$ (upstream)		−0.008
		(0.021)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)		0.033**
		(0.014)
$\pi_{j,t}^{e,h}(\mathrm{up}) \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$		−0.008***
		(0.002)
$T_{j,t} \times \pi_{j,t}^{e,h}(\mathrm{up}) \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$		0.011***
		(0.003)
Downstream (customer) signals		
$g_{j,t}^{r}$	−0.001	−0.001
	(0.001)	(0.001)
$T_{j,t} \times g_{j,t}^{r}$	0.002	0.002
	(0.002)	(0.002)
$g_{j,t}^{r} \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$	0.000	0.000
	(0.000)	(0.000)
$T_{j,t} \times g_{j,t}^{r} \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$	−0.000	−0.000
	(0.001)	(0.001)
$\pi_{j,t}^{e,h}$ (downstream)		−0.025
		(0.023)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)		0.071***
		(0.018)
$\pi_{j,t}^{e,h}(\mathrm{dn}) \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$		−0.002
		(0.002)
$T_{j,t} \times \pi_{j,t}^{e,h}(\mathrm{dn}) \times \widetilde{\ell}_{ij}^{(j),\mathrm{val}}$		0.003
		(0.003)
Lagged expectation $\pi_{i,t-1}^{e,h}$	0.249***	0.248***
	(0.061)	(0.061)
Observations	166,370	166,370
R^2	0.065	0.066

Notes: Pooled-direction specification with linear interactions in centered log partner-side value-based share (cumulated over the full sample). Column (1) includes only outcome signals (supplier price growth upstream, customer purchase growth downstream) and their treatment and size interactions. Column (2) adds partner beliefs and their interactions. All specifications use uniform within-firm link weights, restrict to untreated focal firms ($T_{i,t} = 0$), and pool horizons $h \in \{6, 12\}$. Direction-specific fixed effects, buyer fixed effects, and the standard baseline controls are included throughout. Standard errors use the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback when the variance matrix is not positive definite.

O Relationship persistence: partner-side full specification

This appendix reports the full partner-side persistence specification referenced in Section 4. Table 4 column (2) reports the $T_{j,t}$-interaction terms; the table below shows all estimated coefficients. Own-side persistence results are reported in Table M.2 column (2).

Table O.1: Pooled-direction persistence interactions: partner-side share $\widetilde{\ell}_{ij}^{(j),\text{per}}$

	(1) Outcome only	(2) Outcome + beliefs
Common heterogeneity controls		
$\widetilde{\ell}_{ij}^{(j),\text{per}}$	−0.029*	0.015
	(0.015)	(0.015)
$T_{j,t} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$	0.058*	−0.010
	(0.028)	(0.014)
Upstream (supplier) signals		
$g_{j,t}^{p}$	0.002	0.003
	(0.004)	(0.004)
$T_{j,t} \times g_{j,t}^{p}$	0.002	0.001
	(0.006)	(0.007)
$g_{j,t}^{p} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$	0.002	0.003*
	(0.002)	(0.002)
$T_{j,t} \times g_{j,t}^{p} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$	−0.005	−0.011**
	(0.005)	(0.005)
$\pi_{j,t}^{e,h}$ (upstream)		−0.012
		(0.022)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (upstream)		0.042***
		(0.014)
$\pi_{j,t}^{e,h}(\text{up}) \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		−0.022***
		(0.004)
$T_{j,t} \times \pi_{j,t}^{e,h}(\text{up}) \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		0.035***
		(0.006)
Downstream (customer) signals		
$g_{j,t}^{r}$	−0.001	−0.001
	(0.001)	(0.001)
$T_{j,t} \times g_{j,t}^{r}$	0.002	0.002
	(0.002)	(0.002)
$g_{j,t}^{r} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$	0.000	0.000
	(0.001)	(0.000)
$T_{j,t} \times g_{j,t}^{r} \times \widetilde{\ell}_{ij}^{(j),\text{per}}$	−0.001	−0.000
	(0.001)	(0.001)
$\pi_{j,t}^{e,h}$ (downstream)		−0.024
		(0.026)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)		0.075***
		(0.019)
$\pi_{j,t}^{e,h}(\text{dn}) \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		−0.008
		(0.007)
$T_{j,t} \times \pi_{j,t}^{e,h}(\text{dn}) \times \widetilde{\ell}_{ij}^{(j),\text{per}}$		0.003
		(0.007)
Lagged expectation $\pi_{i,t-1}^{e,h}$	0.249***	0.248***
	(0.061)	(0.061)
Observations	166,370	166,370
R^2	0.065	0.067

Notes: Pooled-direction specification with linear interactions in centered log partner-side persistence share. Estimation details as in Table N.1.

P Nonparametric bin regressions

This appendix provides nonparametric checks by estimating the pooled-direction regression separately within terciles of size or breadth. Each table reports partner-side (Panel A) and own-side (Panel B) definitions. Bins are defined from the untreated-buyer sample prior to stacking.

Value-based size terciles. Table P.1 estimates the regression separately by terciles of value-based size. For partner-side size (Panel A), the treated-partner belief premium on upstream expectations is concentrated in the small-share bin (0.082, significant at 1%), with near-zero estimates in the medium and large bins. This nonmonotonic pattern stands in apparent tension with the positive linear interaction in Table 4 column (1). The resolution lies in the within-bin slopes: the triple interaction $T_{j,t} \times \pi^{e,h}_{j,t}(\text{up}) \times \log \text{Size}$ is positive and significant within the small bin (0.020), indicating that even within the bottom tercile the premium increases steeply with size. Under the own-side share (Panel B), the upstream belief premium is positive and significant in the small (0.083) and medium (0.046) bins.

Table P.1: Pooled-direction bin regressions: value-based size terciles

	Panel A: partner-side share $\widetilde{\ell}_{ij}^{(j),\text{val}}$			Panel B: own-side share $\widetilde{\ell}_{ij}^{(i),\text{val}}$		
	(1) Small	(2) Medium	(3) Large	(4) Small	(5) Medium	(6) Large
Upstream (supplier) signals						
$g^p_{j,t}$	0.010	−0.004	0.012	0.003	0.009*	−0.016
	(0.007)	(0.007)	(0.016)	(0.019)	(0.005)	(0.023)
$T_{j,t} \times g^p_{j,t}$	−0.004	0.000	−0.022	−0.002	−0.003	0.003
	(0.011)	(0.008)	(0.017)	(0.020)	(0.006)	(0.024)
$\pi^{e,h}_{j,t}$ (upstream)	−0.080**	0.017	0.052*	−0.026	−0.010	0.032
	(0.037)	(0.021)	(0.029)	(0.025)	(0.020)	(0.042)
$T_{j,t} \times \pi^{e,h}_{j,t}$ (upstream)	0.082***	0.014	0.006	0.083***	0.046***	0.090
	(0.025)	(0.025)	(0.061)	(0.027)	(0.015)	(0.062)
Downstream (customer) signals						
$g^r_{j,t}$	−0.008*	−0.002	0.001	−0.009*	−0.003*	−0.003
	(0.004)	(0.002)	(0.004)	(0.005)	(0.002)	(0.004)
$T_{j,t} \times g^r_{j,t}$	0.007	0.002	0.000	0.013**	0.002	0.007
	(0.005)	(0.002)	(0.006)	(0.006)	(0.003)	(0.005)
$\pi^{e,h}_{j,t}$ (downstream)	−0.142**	−0.009	0.044	−0.023	0.017	0.041
	(0.059)	(0.021)	(0.055)	(0.048)	(0.015)	(0.043)
$T_{j,t} \times \pi^{e,h}_{j,t}$ (downstream)	0.164***	0.053	−0.007	0.083*	0.030**	−0.022
	(0.050)	(0.032)	(0.065)	(0.048)	(0.013)	(0.046)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.318***	0.297***	0.298***	0.221***	0.291***	0.250***
	(0.046)	(0.032)	(0.036)	(0.055)	(0.039)	(0.076)
Observations	55,986	54,672	55,712	55,468	56,194	54,708
R^2	0.077	0.097	0.094	0.070	0.093	0.064

Notes: Pooled-direction specification estimated separately within terciles of value-based relationship size. Panel A bins by partner-side share; Panel B by own-side share. Bins are defined from the untreated-buyer sample prior to stacking. Each column retains the full set of direction-specific signal interactions including within-bin linear size slopes and size interactions, which are estimated but suppressed for brevity. Estimation details as in Table N.1.

Breadth terciles. Table P.2 confirms the breadth patterns nonparametrically. Under partner breadth (Panel A), the upstream belief premium $T_{j,t} \times \pi^{e,h}_{j,t}(\text{up})$ is positive and significant in the narrow-breadth bin (0.082) and the medium-breadth bin (0.059), but flips sign in the broad-breadth bin (−0.317), though the latter is estimated with substantial noise (s.e. 0.112). The downstream belief premium shows the opposite pattern: it is insignificant in the narrow bin (−0.020) but large and significant in the broad bin (0.177, s.e. 0.062).

Under own breadth (Panel B), the upstream belief premium is significant in the narrow bin (0.071, s.e. 0.039) but small and insignificant in the medium and broad bins. The downstream belief premium is significant in the narrow bin (0.038) and in the broad bin (0.015).

Table P.2: Pooled-direction bin regressions: breadth terciles

	Panel A: partner breadth $\widetilde{b}^{(j)}_{j,t}$			Panel B: own breadth $\widetilde{b}^{(i)}_{i,t}$		
	(1) Narrow	(2) Medium	(3) Broad	(4) Narrow	(5) Medium	(6) Broad
Upstream (supplier) signals						
$g^{p}_{j,t}$	−0.025	0.007	0.081	0.011	−0.001	0.021
	(0.030)	(0.006)	(0.050)	(0.010)	(0.004)	(0.025)
$T_{j,t} \times g^{p}_{j,t}$	0.012	−0.002	−0.026	−0.031	−0.004	0.011
	(0.032)	(0.006)	(0.036)	(0.023)	(0.004)	(0.011)
$\pi^{e,h}_{j,t}$ (upstream)	−0.035	−0.022	0.213**	0.076*	−0.013	0.145
	(0.033)	(0.019)	(0.080)	(0.042)	(0.019)	(0.131)
$T_{j,t} \times \pi^{e,h}_{j,t}$ (upstream)	0.082***	0.059***	−0.317**	0.076*	0.019	−0.013
	(0.028)	(0.019)	(0.112)	(0.044)	(0.012)	(0.017)
Downstream (customer) signals						
$g^{r}_{j,t}$	−0.001	0.002	0.000	−0.001	−0.002	−0.001
	(0.002)	(0.002)	(0.005)	(0.005)	(0.001)	(0.001)
$T_{j,t} \times g^{r}_{j,t}$	0.001	−0.002	−0.000	0.008	0.005**	0.002**
	(0.002)	(0.002)	(0.006)	(0.008)	(0.002)	(0.001)
$\pi^{e,h}_{j,t}$ (downstream)	0.023	−0.009	−0.143*	−0.009	0.044**	0.025**
	(0.014)	(0.025)	(0.078)	(0.030)	(0.015)	(0.011)
$T_{j,t} \times \pi^{e,h}_{j,t}$ (downstream)	−0.020	0.029	0.177***	0.034*	−0.002	0.014**
	(0.019)	(0.022)	(0.062)	(0.019)	(0.015)	(0.006)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.311***	0.299***	0.319***	0.223***	0.353***	0.524***
	(0.026)	(0.042)	(0.041)	(0.066)	(0.043)	(0.066)
Observations	53,350	57,468	55,328	54,848	57,860	53,662
R^2	0.109	0.097	0.078	0.055	0.158	0.306

Notes: Pooled-direction specification estimated separately within terciles of breadth. Panel A bins by partner breadth; Panel B by own breadth. Bins are defined from the untreated-buyer sample prior to stacking. Each column retains the full set of direction-specific signal interactions including within-bin linear breadth slopes and breadth interactions, which are estimated but suppressed for brevity. Estimation details as in Table N.1.

Q Learning from buyers: heterogeneity

This appendix reports additional evidence on downstream information transmission, that is, learning by sellers from their customers. The empirical design mirrors the supplier-learning

analysis in the main text, but restricts attention to customer links. We use uniform within-firm link weights and pool horizons $h \in \{6, 12\}$. As in the main analysis, we focus on untreated sellers ($T_{i,t} = 0$), so coefficients can be interpreted as differential learning from treated versus untreated buyers.

Overview. Across specifications, downstream transmission operates primarily through the belief channel: untreated sellers respond more strongly to the inflation expectations of treated buyers than to those of untreated buyers. In contrast, outcome-based learning from buyers is weak and typically imprecisely estimated. Relative to upstream learning from suppliers, heterogeneity in downstream transmission is less systematic across both relationship importance and network exposure.

The main text documents that the downstream belief premium is positive and significant on average (0.071–0.077 across specifications) but exhibits no significant size or breadth gradient under partner-side definitions (Table 5). The results below provide additional detail.

Size-based heterogeneity (customers). The downstream belief premium does not display a clear monotonic relationship with partner-side size. The Δ(P90–P10) for the downstream belief premium is 0.027 (s.e. 0.029, $p = 0.35$) under partner-side value-based size and 0.010 (s.e. 0.027, $p = 0.71$) under partner-side persistence. These estimates are reported in Table 5 (Panels A and B).

Under own-side size, the downstream belief gradient is modestly positive (0.054, $p = 0.17$ for value-based; 0.056, $p = 0.22$ for persistence), as reported in Table M.1 (Panels A and B).

Breadth-based heterogeneity (customers). Under partner breadth, the downstream belief premium is positive and significant at both P10 (0.085) and P90 (0.071) with no significant gradient ($\Delta = -0.014$, $p = 0.20$; Table 5, Panel C). Under own breadth, the downstream belief premium is significant at P10 (0.085, $p < 0.01$) but insignificant at P90 (0.015), with a marginally significant negative gradient ($\Delta = -0.070$, $p = 0.071$; Table M.1, Panel C).

R Additional results for model validation

This appendix provides full specifications and supplementary heterogeneity results supporting Section 6. We report (i) the complete state-dependent expectation regression, (ii) nonparametric breadth tests, (iii) partner-side size heterogeneity, (iv) own-side heterogeneity specifications for bilateral transaction values, and (v) the equilibrium disclosure rate by relationship persistence.

R.1 State-dependent expectations: full specification

Section 4.1 reports reduced-form interaction results with squared inflation. The full estimating equation is:

$$
\begin{aligned}
\pi^e_{i,t} = \alpha_i + \alpha_t + \rho\,\pi^e_{i,t-1} + X_{ij,t}\gamma \\
&+ \beta_0 g_{j,t} + \beta_1(T_j \times g_{j,t}) + \beta_2(T_i \times g_{j,t}) + \beta_3(T_i \times T_j \times g_{j,t}) \\
&+ \delta_0 \pi^e_{j,t} + \delta_1(T_j \times \pi^e_{j,t}) + \delta_2(T_i \times \pi^e_{j,t}) + \delta_3(T_i \times T_j \times \pi^e_{j,t}) \\
&+ \tilde{\pi}^2_t \times \Big[\beta^\pi_0 g_{j,t} + \beta^\pi_1(T_j \times g_{j,t}) + \beta^\pi_2(T_i \times g_{j,t}) + \beta^\pi_3(T_i \times T_j \times g_{j,t})\Big] \\
&+ \tilde{\pi}^2_t \times \Big[\delta^\pi_0 \pi^e_{j,t} + \delta^\pi_1(T_j \times \pi^e_{j,t}) + \delta^\pi_2(T_i \times \pi^e_{j,t}) + \delta^\pi_3(T_i \times T_j \times \pi^e_{j,t})\Big] + \varepsilon_{ij,t}. \qquad (83)
\end{aligned}
$$

All specifications include firm and quarter fixed effects, sector×quarter, size-class×quarter, area×quarter, and T_j×quarter fixed effects. Standard errors are clustered at the firm level.

Remaining baseline coefficients are reported in Table 3 of the main text.

R.1.1 State-dependent learning: alternative weighting schemes

Table R.1 reports the state-dependent interaction estimates from Table 6 under alternative weighting schemes. The qualitative pattern—positive $\hat{\delta}^\pi_1$ for downstream beliefs, negative $\hat{\delta}^\pi_3$ when both parties are treated—is robust across all four schemes.

Table R.1: State-dependent learning: alternative weighting schemes

	Dep. var.: $\pi^{e,h}_{i,t}$ (p.p.)			
	Untreated buyers		All buyers	
	brd	val	brd	unif
Belief channel $\times\tilde{\pi}^2_t$/sd				
δ^π_1: (up)	0.019	0.020	0.043**	0.051***
	(0.017)	(0.018)	(0.017)	(0.016)
δ^π_1: (down)	0.058***	0.044***	0.037***	0.032**
	(0.014)	(0.015)	(0.013)	(0.013)
δ^π_3: (up)			−0.053***	−0.059***
			(0.015)	(0.015)
δ^π_3: (down)			−0.067***	−0.050***
			(0.018)	(0.013)
Observations	166,370	166,370	497,406	497,406
R^2	0.070	0.071	0.083	0.079

Notes: Same specification as Table 6 under alternative weighting schemes. Columns show breadth-weighted (*brd*) and transaction-value (*val*) results for untreated buyers; breadth-weighted (*brd*) and uniform (*unif*) for all buyers. The shr-weighted untreated results are $\hat{\delta}^\pi_1$(up) = 0.018 (0.012), $\hat{\delta}^\pi_1$(down) = 0.040*** (0.008). $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

R.2 Derivation: bilateral trade response to supplier signals

This subsection derives the structural mapping between the bilateral trade regression coefficients and the model's information-transmission parameters. We show that the CES demand structure

generates testable predictions for both the price channel and the belief channel, and that the horse-race specification in Table 11 cleanly separates them.

Trade value and buyer beliefs. Start from CES demand (13) and write bilateral trade value as

$$V_{k\ell,t} \;=\; p_{k\ell,t} \cdot x_{k\ell,t}(p_{k\ell,t}) \;=\; p_{k\ell,t}^{1-\varepsilon_i}\; \mathbb{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]^{\varepsilon_i - 1}\; a_{ij}\, S_{i,t}, \tag{84}$$

where k denotes the buyer (in sector i), ℓ the supplier (in sector j), $Q_{j,t}$ the supplier-sector cost shifter, and $\mathcal{I}_{k,t}$ the buyer's information set. Taking logs,

$$\log V_{k\ell,t} \;=\; (1-\varepsilon_i)\, \log p_{k\ell,t} \;+\; (\varepsilon_i - 1)\, \log \mathbb{E}[Q_{j,t} \mid \mathcal{I}_{k,t}] \;+\; \log(a_{ij}\, S_{i,t}). \tag{85}$$

Bilateral trade value depends on the supplier's price (which reduces demand through substitution) and the buyer's conditional expectation of the outside option $Q_{j,t}$ (which sustains demand when the buyer believes costs are high industry-wide). Both the price and the belief channels operate through $\log \mathbb{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$: any signal that shifts the buyer's expectation of $Q_{j,t}$ affects bilateral trade.

Two channels of information transmission. The buyer's conditional expectation $\mathbb{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$ responds to two classes of supplier-generated signals. The first is the bilateral price $p_{k\ell,t}$, which on non-disclosure links carries information about the supplier's costs through the price residual (18). The second is the supplier's communicated belief $\hat{q}_{j\ell,t}$, which on disclosure links (or through informal communication on any link) directly informs the buyer about aggregate cost conditions.

Let $\omega_{ij,t} \in [0, 1)$ denote the weight the buyer places on the price residual (the price-extraction weight from (19)), and let $\psi_{ij,t} \geq 0$ denote the weight on the supplier's communicated belief. On disclosure links, $\psi_{ij,t} > 0$ and $\omega_{ij,t} = 0$ by sufficiency of the disclosed belief. On non-disclosure links without informal communication, $\psi_{ij,t} = 0$ and $\omega_{ij,t} > 0$.

Price channel. On non-disclosure links, a unit increase in the log price shifts the buyer's posterior log expectation by $\omega_{ij,t}$:

$$\frac{\partial \log \mathbb{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]}{\partial \log p_{k\ell,t}} \;=\; \omega_{ij,t}.$$

Substituting into (85), the trade-value elasticity with respect to the bilateral price is

$$\frac{\partial \log V_{k\ell,t}}{\partial \log p_{k\ell,t}} \;=\; -(\varepsilon_i - 1)(1 - \omega_{ij,t}). \tag{86}$$

On disclosure links ($\omega_{ij,t} = 0$), the elasticity is $1 - \varepsilon_i < 0$: a price increase reduces trade value because the buyer substitutes toward the outside option. On non-disclosure links with $\omega_{ij,t} > 0$, part of the price increase is attributed to higher fundamentals $Q_{j,t}$, dampening substitution. The treatment differential in the trade-value elasticity is

$$\beta_1^p \;=\; (\varepsilon_i - 1)\, \Delta\omega_{ij} \;>\; 0, \tag{87}$$

where $\Delta\omega_{ij} \equiv \omega_{ij}^{TR} - \omega_{ij}^{UT} > 0$ because treatment reduces the supplier's signal noise and makes the price more informative.

Belief channel. The supplier's communicated belief $\hat{q}_{j\ell,t}$ enters the buyer's information set both on disclosure links (where it is observed directly) and potentially on non-disclosure links (through informal communication during repeated interaction—the margin captured by the belief channel in Section 3.2). In the regression, the supplier's survey expectation $\pi^e_{j,t}$ proxies for $\hat{q}_{j\ell,t}$.

The trade-value response to the supplier's belief is

$$\frac{\partial \log V_{k\ell,t}}{\partial \hat{q}_{j\ell,t}} \;=\; (\varepsilon_i - 1)\,\psi_{ij,t} \;-\; (\varepsilon_i - 1)\,\frac{\partial \log p_{k\ell,t}}{\partial \hat{q}_{j\ell,t}}, \tag{88}$$

where the first term is the demand-sustaining effect (higher $\hat{q}_{j\ell,t}$ raises $\mathbf{E}[Q_{j,t} \mid \mathcal{I}_{k,t}]$, which raises the perceived outside option and dampens substitution) and the second is the price pass-through (the supplier raises its price one-for-one with $\hat{q}_{j\ell,t}$ from (16), reducing demand). The treatment differential is

$$\beta^e_1 \;=\; (\varepsilon_i - 1)\Big(\psi^{TR}_{ij,t} - \psi^{UT}_{ij,t}\Big) \;>\; 0, \tag{89}$$

which is positive because treatment reduces $\sigma^2_{\nu,j}$, making $\hat{q}_{j\ell,t}$ more precise and raising the weight $\psi_{ij,t}$ the buyer places on it.

Equations (87) and (89) are derived holding the disclosure regime fixed. In practice, the treatment may also switch links from non-disclosure to disclosure (the extensive margin of Proposition 2), changing $\omega_{ij,t}$ and $\psi_{ij,t}$ discontinuously. The regression pools over both margins: the estimated $\hat{\beta}^p_1$ and $\hat{\beta}^e_1$ are averages across disclosure regimes, weighted by the share of links in each. Both margins have the same sign for each coefficient, so the sign predictions are unambiguous; the quantitative mapping from regression coefficients to structural parameters requires the regime shares implied by the calibrated disclosure set (Section 6.1).

Horse-race specification. The estimating equation in Section 6.2 includes both channels simultaneously:

$$\begin{aligned}\log V_{k\ell,t} = \alpha_k + \delta_{s(j),t} + \alpha_{T_j,t} + \beta^p_0\, g^p_{j,t} + \beta^p_1\,(T_j \times g^p_{j,t}) \\ + \beta^e_0\, \pi^e_{j,t} + \beta^e_1\,(T_j \times \pi^e_{j,t}) + \varepsilon_{k\ell,t},\end{aligned} \tag{90}$$

where $\delta_{s(j),t}$ are supplier-industry×quarter fixed effects and the remaining notation follows equation (29). The supplier-industry fixed effects absorb industry-level price variation, so β^p_1 is identified from within-industry cross-supplier price differences, while β^e_1 is identified from within-industry cross-supplier differences in communicated beliefs.

The model predicts $\beta^p_1 > 0$ and $\beta^e_1 > 0$, but the two coefficients need not have the same magnitude or empirical robustness. The price channel operates through $\Delta\omega_{ij}$, which depends on the signal-to-noise ratio $\sigma^2_q/\sigma^2_\zeta$ and is attenuated when industry-level price variation is absorbed by fixed effects. The belief channel operates through $\Delta\psi_{ij}$, which depends on the precision gain from treatment and is robust to industry controls because the supplier's expectation $\pi^e_{j,t}$ varies at the supplier level. The empirical finding that $\hat{\beta}^p_1 \approx 0$ while $\hat{\beta}^e_1 > 0$ in the horse-race specification (Table 11) indicates that the allocative consequences of information transmission operate primarily through disclosed beliefs rather than through price-based inference.

R.3 Bilateral transactions: robustness

This appendix reports robustness checks for the bilateral trade results in Table 11. The main finding is that the treated-supplier belief coefficient $\hat{\beta}^e_1 \approx 0.08$–$0.10$ is robust to trimming small

transactions, alternative clustering, persistence restrictions, and all four weighting schemes. The treated-supplier price coefficient $\hat{\beta}_1^p$ is insignificant throughout.

R.3.1 Trimming sensitivity

Table R.2 reports the horse-race specification at five minimum transaction-value thresholds, progressively excluding small links. The belief coefficient is significant at every threshold, with point estimates between 0.08 and 0.11 and t-statistics consistently above 2.9. The coefficient actually strengthens at higher thresholds: at ≥5,000 EUR, $\hat{\beta}_1^e = 0.084^{***}$ ($t = 3.46$), compared to 0.104^{***} ($t = 3.31$) at the ≥1 EUR baseline. The price interaction $\hat{\beta}_1^p$ is insignificant at every threshold.

Table R.2: Bilateral trade belief coefficient: trimming sensitivity

	Dep. var.: $\log \Gamma_{k\ell,t}$				
	(1) ≥1	(2) ≥100	(3) ≥500	(4) ≥1,000	(5) ≥5,000
$g_{j,t}^p$	0.021^{**}	0.028^{***}	0.013	0.017^{**}	0.004
	(0.010)	(0.010)	(0.009)	(0.008)	(0.007)
$T_j \times g_{j,t}^p$	0.016	0.003	0.013	0.002	0.001
	(0.012)	(0.011)	(0.010)	(0.009)	(0.008)
$\pi_{j,t}^e$	-0.109^{***}	-0.102^{***}	-0.096^{***}	-0.089^{***}	-0.077^{***}
	(0.027)	(0.028)	(0.023)	(0.023)	(0.020)
$T_j \times \pi_{j,t}^e$	0.104^{***}	0.113^{***}	0.095^{***}	0.083^{***}	0.084^{***}
	(0.031)	(0.032)	(0.028)	(0.028)	(0.024)
Observations	54,912	49,355	42,208	38,327	27,261
R^2	0.604	0.543	0.524	0.538	0.592

Notes: Each column restricts the sample to links with bilateral trade value above the indicated threshold. Specification and fixed effects as in Table 11, column (1). Standard errors clustered at the buyer firm level. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

By contrast, the price-only specification (without supplier expectations on the right-hand side) produces a treated-supplier price coefficient of $\hat{\beta}_1^p = 0.022^{**}$ ($t = 2.41$) at the ≥1 EUR baseline, but this coefficient collapses to insignificance at ≥100 EUR ($\hat{\beta}_1^p = 0.007$, $t = 0.90$) and is zero or negative at all higher thresholds (Appendix Table R.5). The baseline price-channel result is driven entirely by transactions below 100 EUR—approximately 10% of the sample—where the structural interpretation (buyer learns about supplier fundamentals from a 50-euro invoice) is least plausible. Adding supplier-industry×quarter fixed effects eliminates the price coefficient even at the ≥1 EUR baseline ($\hat{\beta}_1^p = 0.013$, $t = 1.09$).

R.3.2 Clustering

The baseline clusters standard errors at the buyer level, matching the within-buyer cross-supplier identifying variation. Because the supplier's expectations $\pi_{j,t}^e$ vary at the supplier×quarter level, supplier-level clustering is a relevant robustness check. Table R.3 reports the belief coefficient under three clustering schemes. The point estimate is identical ($\hat{\beta}_1^e = 0.104$); the standard error increases from 0.031 (buyer) to 0.037 (supplier) to 0.039 (two-way), with the coefficient remaining significant at the 1% level under all three schemes ($t = 2.67$ under two-way clustering).

The modest SE inflation (26%) contrasts with the much larger inflation seen for the price coefficient (where supplier clustering tripled the SE from 0.9 to 2.9). This is because $\pi^e_{j,t}$ has substantially more within-supplier time-series variation than $g^p_{j,t}$: expectations change quarter to quarter as firms update their views, while price growth is more persistent within supplier.

Table R.3: Bilateral trade belief coefficient: clustering sensitivity

Clustering	$\hat{\beta}^e_1$ on $T_j \times \pi^e_{j,t}$	s.e.	t-stat
Buyer firm	0.104	0.031	3.31
Supplier firm	0.104	0.037	2.82
Buyer + supplier (two-way)	0.104	0.039	2.67

Notes: All rows estimate the horse-race specification from Table 11, column (1), with uniform weights and supplier-industry×quarter fixed effects. Only the variance estimator changes across rows.

R.3.3 Persistence and link duration

Table R.4 reports the belief coefficient on subsamples defined by minimum link duration and minimum transaction value. On links lasting four or more quarters, the coefficient is 0.135^{***} ($t = 3.04$) at the ≥1 EUR threshold and 0.114^{***} ($t = 4.33$) at ≥5,000 EUR—the strongest single estimate in the analysis. On the most persistent links (≥8 quarters), the coefficient remains positive but is imprecisely estimated at lower thresholds due to the smaller sample; at ≥5,000 EUR, it is 0.084^{**} ($t = 2.11$), consistent with the disclosure channel operating on large, stable trade relationships.

Table R.4: Bilateral trade belief coefficient: persistence

	$\hat{\beta}^e_1$ on $T_j \times \pi^e_{j,t}$		
Min. trade value	Full sample	≥4 quarters	≥8 quarters
≥1 EUR	0.104^{***}	0.135^{***}	0.061
	(0.031)	(0.044)	(0.054)
≥1,000 EUR	0.083^{***}	0.099^{***}	0.044
	(0.028)	(0.031)	(0.043)
≥5,000 EUR	0.084^{***}	0.114^{***}	0.084^{**}
	(0.024)	(0.026)	(0.040)

Notes: Each cell reports $\hat{\beta}^e_1$ from the horse-race specification with supplier-industry×quarter fixed effects and uniform weights. Rows restrict by minimum bilateral trade value; columns restrict by minimum link duration (number of quarters the dyad is observed). Standard errors (in parentheses) clustered at the buyer firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

R.3.4 Price-only specification: trimming sensitivity

For completeness, Table R.5 reports the price-only specification ($T_j \times g^p_{j,t}$ without supplier expectations) under both baseline and supplier-industry×quarter fixed effects, at five transaction-value thresholds. The baseline $\hat{\beta}^p_1 = 0.022^{**}$ at ≥1 EUR is driven by transactions below 100 EUR and

vanishes with trimming or industry controls. This fragility motivates the horse-race specification in the main text, which reveals that the robust bilateral trade effect operates through beliefs, not prices.

Table R.5: Price-only specification: trimming and industry controls

	Min. bilateral trade value				
	≥1	≥100	≥500	≥1,000	≥5,000
Panel A: Baseline FE					
$T_j \times g^p_{j,t}$	0.022**	0.007	−0.001	−0.008	−0.011
	(0.009)	(0.008)	(0.007)	(0.007)	(0.008)
Panel B: + Supplier industry × quarter FE					
$T_j \times g^p_{j,t}$	0.013	−0.000	0.010	−0.001	−0.001
	(0.012)	(0.012)	(0.010)	(0.009)	(0.008)
Observations	55,931	50,378	43,247	39,380	28,391

Notes: Each column restricts the sample by minimum bilateral trade value. Panel A includes buyer FE, buyer characteristics×quarter FE, and T_j×quarter FE. Panel B additionally includes supplier-industry×quarter FE. Uniform weights throughout. Standard errors clustered at the buyer firm level. N in Panel B is slightly lower due to singleton observations dropped by industry×quarter FE. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

R.3.5 Supplier-side controls

Table R.6 adds supplier size-class×quarter fixed effects to the baseline specification. These absorb time-varying differences in trade volumes across supplier size categories that might correlate with both expectations and procurement. The belief coefficient *strengthens*: $\hat{\beta}^e_1 = 0.114^{***}$ ($t = 3.82$), compared with 0.104^{***} ($t = 3.31$) in the baseline. The coefficient is robust to trimming at all thresholds under this richer specification ($t = 2.83$ to 3.82). Adding supplier size controls sharpens the result rather than attenuating it, confirming that the cross-sectional correlation is not driven by supplier size composition.

R.3.6 Conditioning on lagged trade

The baseline specification uses log bilateral trade value in levels. A natural concern is that the belief coefficient reflects persistent cross-link differences in trade volumes that correlate with supplier characteristics. To address this, we add lagged log bilateral trade value $\log \Gamma_{k\ell,t-1}$ as a control, conditioning the regression on the link's trade level in the previous quarter. This restricts the sample to links observed in consecutive quarters ($N = 26{,}087$, roughly half the baseline).

Table R.7 reports the results. The lagged dependent variable has a coefficient of 0.81 ($t = 76$), indicating strong persistence. The belief coefficient drops from 0.182 in the restricted sample without the lag to $\hat{\beta}^e_1 = 0.063^{***}$ ($t = 2.86$) with the lag. The attenuation—roughly two-thirds—reflects the fact that much of the cross-link variation is persistent and absorbed by lagged trade. The surviving coefficient captures the incremental effect: even conditional on the previous quarter's trade level, treated suppliers' expectations predict higher contemporaneous trade.

Table R.6: Belief coefficient with supplier size controls

	Dep. var.: $\log \Gamma_{k\ell,t}$				
	(1)	(2)	(3)	(4)	(5)
	Baseline	+ Supplier size × quarter FE			
		≥1	≥500	≥1,000	≥5,000
$T_j \times g^p_{j,t}$	0.016	0.012	0.010	0.002	0.002
	(0.012)	(0.012)	(0.011)	(0.009)	(0.008)
$T_j \times \pi^e_{j,t}$	0.104***	0.114***	0.092***	0.079***	0.071***
	(0.031)	(0.030)	(0.028)	(0.028)	(0.024)
Sup. industry × qtr	Yes	Yes	Yes	Yes	Yes
Sup. size × qtr	No	Yes	Yes	Yes	Yes
Observations	54,912	54,912	42,208	38,327	27,261

Notes: Column (1) reproduces the baseline from Table 11, column (1). Columns (2)–(5) add supplier size-class×quarter fixed effects and progressively trim by minimum bilateral trade value. Uniform weights throughout. Standard errors clustered at the buyer firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The base effect $\hat{\beta}^e_0$ drops from -0.138 to -0.025 and becomes insignificant, confirming that the cross-sectional penalty on untreated suppliers' expectations is largely a level phenomenon absorbed by lagged trade.

R.3.7 Within-link variation: dyad fixed effects and first differences

Under dyad fixed effects, which absorb all time-invariant link-level heterogeneity and identify the coefficient from within-link time-series variation alone, the belief coefficient is $\hat{\beta}^e_1 = 0.003$ ($t = 0.19$) and the R^2 rises to 0.93. The first-difference specification ($\Delta \log \Gamma_{k\ell,t}$ on the left-hand side) yields $\hat{\beta}^e_1 = 0.035$ ($t = 1.67$).

These results are consistent with the disclosure mechanism operating primarily through *persistent* cross-link differences rather than quarter-to-quarter fluctuations. In the model, treatment status is fixed and disclosure is a level decision: a treated supplier is persistently more informative, and the buyer persistently allocates more trade to that supplier relative to untreated suppliers reporting similar expectations. Dyad fixed effects absorb precisely this persistent cross-link variation, which is the object the model targets. The within-dyad time-series variation in $\pi^e_{j,t}$ has a standard deviation of 2.1 percentage points, comparable to the cross-sectional residual (1.3 pp after industry×quarter effects), but it largely reflects quarter-to-quarter noise in survey responses rather than genuine changes in the supplier's informativeness. The lagged-trade specification, which conditions on the previous quarter's trade level while preserving cross-link variation, provides the more informative robustness check and confirms that the belief coefficient survives ($t = 2.86$).

R.4 Equilibrium disclosure by relationship persistence

Figure R.1 plots the equilibrium disclosure rate by decile of relationship persistence ρ_{ij}, restricting to the 2,787 links with observed (non-imputed) ρ_{ij}. The overall gradient is positive: disclosure rises from near zero below $\rho \approx 0.30$ to above 40% around the median. The gradient

Table R.7: Belief coefficient conditioning on lagged trade

	Dep. var.: $\log \Gamma_{k\ell,t}$				
	(1)	(2)	(3)	(4)	(5)
	No lag	+ Lag	+ Lag, by threshold		
			≥500	≥1,000	≥5,000
$\log \Gamma_{k\ell,t-1}$		0.809***	0.822***	0.818***	0.839***
		(0.011)	(0.011)	(0.011)	(0.012)
$T_j \times g^p_{j,t}$	0.008	−0.009	−0.009	−0.009	−0.009
	(0.022)	(0.007)	(0.007)	(0.007)	(0.006)
$\pi^e_{j,t}$	−0.138***	−0.025			
	(0.043)	(0.017)			
$T_j \times \pi^e_{j,t}$	0.182***	0.063***	0.042**	0.029	0.034**
	(0.049)	(0.022)	(0.018)	(0.018)	(0.016)
Observations	26,087	26,087	20,396	18,576	13,413
R^2	0.699	0.895	0.899	0.902	0.913

Notes: Sample restricted to links observed in consecutive quarters. Column (1) reproduces the horse-race specification on the restricted sample without the lag. Column (2) adds $\log \Gamma_{k\ell,t-1}$. Columns (3)–(5) additionally trim by minimum bilateral trade value. All specifications include supplier-industry×quarter FE and uniform weights. Standard errors clustered at the buyer firm level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

is driven by $R_{ij} = \beta\rho_{ij}/(1 - \beta\rho_{ij})$, which amplifies the cooperative surplus on persistent links and relaxes the IC condition.

The pattern is non-monotone: disclosure declines in the right tail, above $\rho \approx 0.50$. This reflects the joint distribution of ρ_{ij} and a_{ij} in the data. Links in the highest ρ bins tend to have below-average expenditure shares ($a_{ij} \approx 0.009$ for $\rho \in [0.6, 0.7]$ versus $a_{ij} \approx 0.022$ for $\rho \in [0.4, 0.5]$), so the rising persistence is more than offset by the falling bilateral trade volume. The IC condition binds on the *product* $R_{ij} \cdot W(\omega_j) \cdot \Gamma_{ij}$, not on ρ_{ij} alone, and the negative correlation between persistence and size in the observed link distribution causes disclosure to turn down at high ρ values.

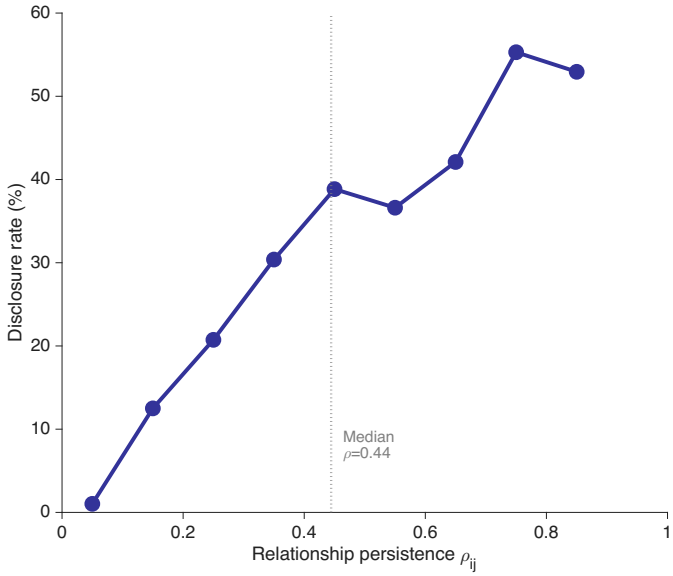


Figure R.1: Equilibrium disclosure rate by relationship persistence

Notes: Each dot reports the disclosure rate among active links in the calibrated model within a decile of relationship persistence ρ_{ij} (quarterly survival probability), restricting to observed (non-imputed) links. The vertical dotted line marks the median $\rho = 0.44$ among links with $\rho > 0$. Disclosure rises from near zero at low persistence to above 40% around the median, then declines at high ρ values because high-persistence links tend to have smaller expenditure shares a_{ij}, more than offsetting the effect of higher R_{ij} on the IC condition. Calibrated parameters as in Table 9.

S Calibration: details and implementation

This appendix provides full details on the calibration procedure, the construction of each targeted moment, the identification mapping from moments to parameters, and the computational algorithm.

S.1 Parameters and identification

The model has two classes of parameters. *External parameters* are set directly from the data without optimization: the demand elasticity $\varepsilon = 6$, the discount factor $\beta = 0.99$, the common-component variance $\hat{\sigma}^2_{q,j}$ (from time-series variance of industry-average price growth), and the private-signal noise $\hat{\sigma}^2_{\nu,j}$ (from the price–cost pass-through regression). *Internal parameters* are (f, λ_ζ), calibrated by minimum distance. The disclosure cost is a uniform setup cost f with no per-period component; see Section 6.1 for justification.

The raw idiosyncratic price dispersion $\hat{\sigma}^2_{\zeta,j}$ (from supplier price dispersion within 2-digit ATECO sectors) is scaled by λ_ζ before entering the model: $\sigma^2_{\zeta,j} = \lambda_\zeta \, \hat{\sigma}^2_{\zeta,j}$.

Signal parameters. The private-signal noise $\hat{\sigma}^2_{\nu,j}$ is identified from the industry-level price–cost pass-through regression (27). The coefficient $\hat{\alpha}_j$ estimates the Kalman gain $\alpha_j = \sigma^2_q/(\sigma^2_q + \sigma^2_\nu)$, from which $\hat{\sigma}^2_{\nu,j} = \hat{\sigma}^2_{q,j}(1 - \hat{\alpha}_j)/\hat{\alpha}_j$. This approach identifies σ^2_ν from the pricing equation directly, without relying on survey-based inflation expectations. Industries without a valid pass-through estimate (negative or above 1) are imputed with the cross-industry median ($\hat{\alpha} = 0.48$). For the treatment simulation, we scale the baseline $\hat{\sigma}^2_{\nu,j}$ by the ratio of treated-to-untreated

within-cell forecast dispersion (median ratio 0.77), capturing the treatment's effect on the buyer's demand-side information.

The industry-level cost volatility $\hat{\sigma}^2_{\zeta,j}$ is estimated from supplier price index dispersion across 2-digit ATECO sectors in the bilateral transaction data. The common-component variance $\sigma^2_{q,j}$ is estimated from the time-series variance of industry-level output price inflation.

Disclosure setup cost. f is the per-link cost of establishing a verifiable disclosure protocol. It is identified by the upstream belief premium (moment m1): a higher f raises the IC threshold, reducing the equilibrium disclosure rate and hence the average weight buyers place on treated suppliers' beliefs.

Price-noise scaling. λ_ζ adjusts the effective idiosyncratic price noise relative to the raw administrative estimate. It is identified by the upstream price coefficient (moment m2): a higher λ_ζ makes prices noisier, reducing the price-signal weight $\omega_{p,j}$ and hence the treatment-induced shift in price-based learning. The scaling corrects for demand shocks, product-mix variation, and measurement error in the raw within-industry price dispersion that are orthogonal to the model's idiosyncratic cost component.

S.2 Targeted moments and their empirical counterparts

All moments are computed from the treated subsample of the SIGE experiment (64 of 78 industries have $\Delta\tau_j > 0$). Let $T_j \in \{0, 1\}$ denote the treatment indicator, $\pi^e_{i,t}$ the one-quarter-ahead price expectation of firm i in industry j, and $g^p_{j,t}$ the supplier price growth of firm i's main supplier from the bilateral transaction data.

Moment m1: upstream belief premium.

$$\hat{\beta}_1^{\text{up}} = \text{Cov}(T_j,\, \pi^e_{i,t})/\text{Var}(T_j), \tag{91}$$

estimated by OLS with industry and quarter fixed effects, clustering by supplier industry. The model analog is $\text{mean}_{i\in\text{treat}}(\Delta\psi_{ij,t})$ averaged over all upstream links (i, j) in treated industries, where $\Delta\psi_{ij,t}$ is the treatment-induced increase in the Kalman weight on the supplier's belief.

Moment m2: upstream price coefficient.

$$\hat{\beta}_2 = \text{Cov}(T_j \cdot g^p_{j,t},\, \pi^e_{i,t})/\text{Var}(T_j \cdot g^p_{j,t}), \tag{92}$$

estimated by OLS with buyer–quarter fixed effects. This measures how much price signals from treated suppliers pass into buyer beliefs. The model analog is the average treatment-induced shift in the price-signal weight $\Delta\omega_{ij}$ on upstream links, which depends on λ_ζ through $\omega_{p,j} = \sigma^2_{q,j}/(\sigma^2_{q,j} + \lambda_\zeta\hat{\sigma}^2_{\zeta,j})$.

S.3 Objective function

The calibration minimizes a diagonal GMM criterion over the two targeted moments:

$$\mathcal{L}(\theta) = \sum_{k=1}^{2} \left(\frac{m_k(\theta) - \hat{m}_k}{\hat{s}_k} \right)^2, \tag{93}$$

where $m_k(\theta)$ is the model-implied analog of moment k at parameters $\theta = (f, \lambda_\zeta)$, $\hat{m}_k$ is the empirical estimate, and $\hat{s}_k$ is its standard error. The system is exactly identified: two parameters, two moments. Standard errors are computed with two-way clustering by buyer and quarter for belief moments, and supplier clustering for price moments.

S.4 Computational algorithm

The calibration proceeds in two stages.

Stage 1: coarse grid search. We evaluate $\mathcal{L}(\theta)$ on a $10 \times 9 = 90$-point grid spanning $\log_{10}(f) \in [-12, -4]$ and $\log_{10}(\lambda_\zeta) \in [-0.5, 1.0]$. At each grid point we solve the full equilibrium fixed point and compute both model moments. The grid search identifies the global basin of attraction and verifies that the objective surface is unimodal over the relevant range.

Stage 2: local refinement. Starting from the grid optimum, we apply `fminsearch` (Nelder–Mead) with 20 perturbation restarts from jittered versions of the Stage 1 optimum (uniform ±15–30% perturbations on each parameter in log space), retaining the global minimum. The exactly-identified system converges to a near-zero objective ($\mathcal{L}^* < 0.01$) in all restarts, confirming global identification.

Welfare formula verification. For each candidate calibration, we verify the welfare decomposition formula against a Monte Carlo simulation with 100,000 draws. The criterion requires that the total consumption-equivalent welfare cost implied by the T1+T2+T3 formula lie within 5% of the simulation-based estimate. All reported calibrations pass this check.

S.5 Sensitivity to signal parameter measurement

The calibrated $\lambda_\zeta = 3.16$ scales the raw within-industry price dispersion to correct for non-cost components. To assess sensitivity, we re-run the calibration under 30 alternative specifications of $\hat{\sigma}^2_{\zeta,j}$, varying (i) the minimum cell size threshold (5, 10, 20, or 50 observations per industry-quarter cell) and (ii) an additive measurement floor $\underline{\sigma}^2 \in \{1, 2, 4, 8, 16\}$. Each specification yields a new set of industry-level $\hat{\sigma}^2_{\zeta,j}$ values, from which we re-run the full calibration and compute the equilibrium markup cost T1.

Across all specifications, the equilibrium T1 ranges from 0.26% to 0.34% of GDP. The baseline estimate of 0.29% lies near the middle of this range, confirming that the welfare result is not sensitive to the precise treatment of measurement error in the raw price dispersion data.

S.6 GDP normalisation and cost-to-surplus ratios

The setup cost f is expressed as a share of GDP and therefore appears numerically small in absolute terms. This subsection shows that the economically relevant object—the cost-to-surplus ratio $f/(R_{ij} \cdot W(\omega_{p,j}) \cdot \Gamma_{ij})$—takes intuitive values: it substantially exceeds one for non-disclosing links and falls below one for disclosing links, consistent with the IC condition.

GDP normalisation. The demand scale $\Gamma_{k\ell,t}$ is defined as buyer k's real expenditure on inputs from supplier ℓ, measured as a share of aggregate value added (GDP). In the calibrated network, the median active link has an expenditure share $a_{ij} \approx 0.001$ and the aggregate Domar weight of the supplying industry is of order 10^{-2}, so the bilateral demand scale satisfies $\Gamma_{k\ell} \sim 10^{-5}$ at the link level. The cost parameter f inherits these units and is therefore small in absolute terms by construction.

Cost-to-surplus ratios. The economically meaningful quantity is the ratio $f/(R_{ij} \cdot W(\omega_{p,j}) \cdot \Gamma_{ij})$, which equals one at the IC threshold. At the median non-disclosing link, the setup cost substantially exceeds the present value of cooperative surplus, placing the link well inside the non-disclosure region. At the median disclosing link, the cooperative surplus exceeds the setup cost, making disclosure privately optimal. The large gap between these groups reflects the strong heterogeneity in R_{ij} (driven by ρ_{ij}) and Γ_{ij} (driven by a_{ij}) across the link distribution.

Interpretation. The absolute magnitude of f is a pure unit artefact of GDP normalisation. The economically meaningful object is the ratio $f/(R_{ij} \cdot W(\omega_j) \cdot \Gamma_{ij})$, which falls as ρ_{ij} rises (through R_{ij}) and as a_{ij} rises (through Γ_{ij}). This generates the disclosure gradients in both relationship persistence and expenditure share visible in Figures R.1 and 5.

T Fixed-effect robustness

This appendix examines the sensitivity of the baseline learning results (Table 3) to alternative fixed-effect structures. Table L.22 (reported above) adds partner-level controls to the main specification. Tables T.1 and T.2 consider two further variants: a simplified fixed-effect structure that drops the treatment-pair×quarter effects, and a specification that combines the simplified structure with partner-level controls.

Simplified fixed effects. Table T.1 replaces the treatment-pair×quarter effects with simple quarter fixed effects. The untreated-only specifications (columns 1–4) remain similar to the baseline: the downstream belief premium δ_1 ranges from 0.032 to 0.056, compared with 0.064 to 0.113 in the main table. The full-sample columns (5–8) lose precision, reflecting the absence of treatment-pair×quarter effects that absorb composition differences across (T_i, T_j) cells—a role that is redundant in the untreated-only sample where $T_i = 0$ throughout.

Table T.1: Robustness: simplified fixed effects

	Untreated only ($T_{i,t}=0$)				Full sample			
	(1) brd	(2) unif	(3) val	(4) shr	(5) brd	(6) unif	(7) val	(8) shr
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^p_{j,t}$)								
β_0: $g^p_{j,t}$	−0.001	0.001	0.001	0.008	−0.022***	−0.018***	−0.019***	−0.011
	(0.003)	(0.004)	(0.004)	(0.008)	(0.007)	(0.006)	(0.006)	(0.009)
β_1: $T_{j,t}\, g^p_{j,t}$	0.006	0.006	−0.004	−0.005	0.018*	0.017*	0.008	0.006
	(0.005)	(0.007)	(0.004)	(0.006)	(0.009)	(0.009)	(0.007)	(0.006)
δ_0: $\pi^{e,h}_{j,t}$	0.025*	0.019	0.008	0.020	−0.077***	−0.077***	−0.077***	−0.063**
	(0.013)	(0.014)	(0.023)	(0.026)	(0.019)	(0.021)	(0.022)	(0.029)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	−0.008	−0.014	0.016	0.016	−0.025	−0.033**	−0.005	−0.016
	(0.013)	(0.011)	(0.018)	(0.019)	(0.016)	(0.014)	(0.018)	(0.023)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					0.028	0.037*	0.008	0.017
					(0.023)	(0.020)	(0.024)	(0.031)
Panel B: Downstream signals ($g_{j,t}$ = customer purchase growth $g^r_{j,t}$)								
β_0: $g^r_{j,t}$	−0.000	−0.000	0.001	0.000	−0.001	−0.001	−0.000	−0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
β_1: $T_{j,t}\, g^r_{j,t}$	0.000	0.000	−0.001	0.002	0.000	0.001	−0.000	0.001
	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
δ_0: $\pi^{e,h}_{j,t}$	−0.005	−0.005	−0.019	0.001	−0.111***	−0.108***	−0.115***	−0.093**
	(0.017)	(0.020)	(0.023)	(0.023)	(0.035)	(0.036)	(0.036)	(0.037)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.041***	0.032***	0.056***	0.045***	0.012	0.010	0.030	0.011
	(0.010)	(0.010)	(0.013)	(0.016)	(0.025)	(0.025)	(0.025)	(0.026)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					−0.018	−0.011	−0.019	−0.002
					(0.034)	(0.035)	(0.032)	(0.031)
ρ: $\pi^{e,h}_{i,t-1}$	0.261***	0.248***	0.261***	0.260***	0.293***	0.284***	0.292***	0.294***
	(0.053)	(0.061)	(0.053)	(0.053)	(0.043)	(0.044)	(0.043)	(0.043)
Observations	166,370	166,370	166,370	166,370	497,406	497,406	497,406	497,406
R^2	0.068	0.061	0.069	0.069	0.106	0.101	0.108	0.105

Notes: Same specification as Table 3, except that the treatment-pair×quarter fixed effects (T_j×quarter in the untreated-only sample; direction×T_i × T_j×quarter in the full sample) are replaced by simple quarter fixed effects. For brevity, β_2, β_3 (outcome channel, own treatment), and δ_2 (belief channel, own treatment) are estimated but suppressed. All other details as in Table 3. $^*p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Simplified fixed effects with partner controls. Table T.2 combines both modifications: simplified quarter effects plus partner-level controls. This is the most conservative specification. The pattern mirrors Table T.1: the untreated-only downstream belief premium δ_1 ranges from 0.033 to 0.052, while full-sample columns lose precision for the same mechanical reason.

Table T.2: Robustness: simplified fixed effects with partner-level controls

	Untreated only ($T_{i,t}=0$)				Full sample			
	(1) brd	(2) unif	(3) val	(4) shr	(5) brd	(6) unif	(7) val	(8) shr
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^p_{j,t}$)								
β_0: $g^p_{j,t}$	−0.007*	−0.002	−0.001	0.005	−0.029***	−0.022***	−0.023***	−0.013
	(0.003)	(0.004)	(0.004)	(0.008)	(0.008)	(0.007)	(0.007)	(0.009)
β_1: $T_{j,t}\, g^p_{j,t}$	0.012*	0.008	−0.001	−0.002	0.025**	0.021**	0.012	0.008
	(0.006)	(0.008)	(0.005)	(0.006)	(0.009)	(0.010)	(0.007)	(0.006)
δ_0: $\pi^{e,h}_{j,t}$	0.042***	0.022	0.011	0.019	−0.064***	−0.071***	−0.068***	−0.060**
	(0.012)	(0.015)	(0.023)	(0.025)	(0.018)	(0.020)	(0.022)	(0.029)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	−0.016	−0.015	0.012	0.009	−0.033**	−0.036**	−0.010	−0.019
	(0.012)	(0.012)	(0.018)	(0.019)	(0.015)	(0.014)	(0.018)	(0.022)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					0.027	0.036*	0.005	0.017
					(0.023)	(0.020)	(0.025)	(0.032)
Panel B: Downstream signals ($g_{j,t}$ = customer purchase growth $g^r_{j,t}$)								
β_0: $g^r_{j,t}$	−0.000	−0.001	−0.000	−0.000	−0.001	−0.001	−0.000	−0.000
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)
β_1: $T_{j,t}\, g^r_{j,t}$	0.000	0.001	−0.000	0.002*	0.000	0.001	−0.001	0.002
	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
δ_0: $\pi^{e,h}_{j,t}$	0.004	−0.009	−0.012	0.013	−0.101***	−0.104**	−0.106***	−0.087**
	(0.019)	(0.023)	(0.025)	(0.024)	(0.033)	(0.036)	(0.035)	(0.036)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.042***	0.035**	0.052***	0.033*	0.010	0.010	0.028	0.009
	(0.012)	(0.012)	(0.015)	(0.016)	(0.024)	(0.025)	(0.024)	(0.026)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					−0.015	−0.009	−0.019	−0.000
					(0.034)	(0.036)	(0.033)	(0.033)
ρ: $\pi^{e,h}_{i,t-1}$	0.265***	0.252***	0.260***	0.260***	0.292***	0.284***	0.292***	0.294***
	(0.050)	(0.057)	(0.051)	(0.052)	(0.043)	(0.044)	(0.043)	(0.043)
Observations	166,370	166,370	166,370	166,370	497,406	497,406	497,406	497,406
R^2	0.070	0.063	0.067	0.068	0.106	0.101	0.108	0.105

Notes: Same specification as Table 3, except that the treatment-pair×quarter fixed effects are replaced by simple quarter fixed effects, and partner-level fixed effects (j's sector×quarter, size class×quarter, and geographic area×quarter) are added. For brevity, β_2, β_3 (outcome channel, own treatment), and δ_2 (belief channel, own treatment) are estimated but suppressed. All other details as in Table 3. $^*p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

U Robustness: excluding reassigned firms

A small number of firms are reassigned across treatment and control groups over the sample period. The baseline specification captures these reassignments through the time-varying treatment indicator $T_{i,t}$. To the extent that reassigned firms retain residual informational gains from prior exposure, the affected observations attenuate the estimated treatment contrast. Table U.1

replicates the baseline learning specification of Table 3 after dropping all firms that switch treatment status at any point during the sample, together with their trading partners. All main results are preserved.

Table U.1: Robustness: excluding reassigned firms and their trading partners

	Signal	Horse race (UNT only)				Triple interactions (ALL)			
	(1) brd	(2) brd	(3) unif	(4) val	(5) shr	(6) brd	(7) unif	(8) val	(9) shr
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^p_{j,t}$)									
β_0: $g^p_{j,t}$	0.004	0.002	0.012*	0.005	0.007	0.001	0.012**	0.006	0.010
	(0.004)	(0.004)	(0.007)	(0.006)	(0.008)	(0.005)	(0.005)	(0.007)	(0.009)
β_1: $T_{j,t}\, g^p_{j,t}$	−0.003	0.003	−0.006	−0.007	−0.007	0.007	−0.005	−0.005	−0.008
	(0.006)	(0.005)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)
β_2: $T_{i,t}\, g^p_{j,t}$						−0.009*	−0.018**	−0.008	−0.011*
						(0.005)	(0.006)	(0.006)	(0.006)
β_3: $T_{i,t}\, T_{j,t}\, g^p_{j,t}$						−0.000	0.015	0.015**	0.014**
						(0.006)	(0.009)	(0.007)	(0.006)
δ_0: $\pi^{e,h}_{j,t}$		−0.033	−0.030	−0.048	−0.007	−0.023	−0.035	−0.047	−0.000
		(0.036)	(0.036)	(0.040)	(0.034)	(0.041)	(0.043)	(0.042)	(0.038)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$		0.073*	0.055*	0.095**	0.057	0.063	0.060*	0.098**	0.060
		(0.040)	(0.029)	(0.038)	(0.035)	(0.043)	(0.032)	(0.043)	(0.038)
δ_2: $T_{i,t}\, \pi^{e,h}_{j,t}$						0.056	0.059	0.068	0.012
						(0.045)	(0.047)	(0.040)	(0.044)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$						−0.052	−0.055*	−0.067**	−0.028
						(0.040)	(0.031)	(0.032)	(0.034)
Panel B: Downstream signals ($g_{j,t}$ = customer purchase growth $g^r_{j,t}$)									
β_0: $g^r_{j,t}$	−0.001	−0.001	−0.002	−0.000	−0.001	0.000	0.000	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)
β_1: $T_{j,t}\, g^r_{j,t}$	0.002	0.002	0.002	0.001	0.003**	−0.000	−0.000	−0.001	0.001
	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)
β_2: $T_{i,t}\, g^r_{j,t}$						0.001	0.001	−0.000	−0.000
						(0.002)	(0.001)	(0.002)	(0.002)
β_3: $T_{i,t}\, T_{j,t}\, g^r_{j,t}$						0.002	0.001	0.003	−0.000
						(0.003)	(0.002)	(0.003)	(0.003)
δ_0: $\pi^{e,h}_{j,t}$		−0.045	−0.050	−0.055	−0.015	−0.059**	−0.042	−0.044	−0.016
		(0.032)	(0.038)	(0.039)	(0.035)	(0.027)	(0.026)	(0.031)	(0.037)
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$		0.108***	0.101***	0.122***	0.080**	0.125***	0.092**	0.099***	0.067
		(0.033)	(0.034)	(0.036)	(0.036)	(0.032)	(0.037)	(0.032)	(0.045)
δ_2: $T_{i,t}\, \pi^{e,h}_{j,t}$						0.117***	0.088***	0.099**	0.059
						(0.040)	(0.030)	(0.043)	(0.047)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$						−0.148***	−0.112**	−0.101**	−0.080
						(0.045)	(0.042)	(0.044)	(0.050)
ρ: $\pi^{e,h}_{i,t-1}$	0.256***	0.255***	0.237***	0.255***	0.254***	0.282***	0.275***	0.282***	0.281***
	(0.044)	(0.044)	(0.053)	(0.044)	(0.044)	(0.044)	(0.045)	(0.043)	(0.044)
Observations	108,154	108,154	108,154	108,154	108,154	400,016	400,016	400,016	400,016
R^2	0.066	0.068	0.058	0.069	0.068	0.080	0.075	0.081	0.080

Notes: Same specification as Table 3, estimated after dropping all firms that are reassigned across treatment and control groups at any point during the sample, together with their direct trading partners. Column (1) includes price signals only; columns (2)–(5) add belief signals (untreated firms only, $T_{i,t} = 0$); columns (6)–(9) report the full-sample specification with own-treatment interactions. Panels A and B are extracted from a single joint regression; ρ and fit statistics are common across directions. Driscoll–Kraay standard errors (bandwidth 2) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.