Disclosure and Inference in Production Networks: Prices and Expectations

Vincenzo Cuciniello, Bank of Italy

Luigi Paciello*, EIEF & CEPR

March 26, 2026

Abstract

When firms learn about common cost pressures from suppliers, they may do so through direct disclosure or by inferring costs from prices. These two modes of learning have different implications for markups, trade, and welfare. Combining a persistent information treatment that shifts firms' inflation beliefs 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 sharply across relationships.

Belief-based learning from suppliers is strongest in durable and economically important relationships, whereas price-based learning is strongest when suppliers source from broader upstream networks. This pattern is driven by supplier-side relationship characteristics—persistence, economic importance, and supplier network size—rather than buyer-side characteristics, pointing to supplier disclosure incentives as the key margin. Consistent with this interpretation, suppliers' beliefs robustly predict bilateral trade volumes.

We develop a model in which suppliers strategically disclose information about costs to buyers, eliminating the markup distortion that arises when buyers otherwise infer costs from prices. Disclosure is sustained by relationship surplus and therefore concentrated in durable, high-surplus relationships, consistent with the heterogeneity we document in the data. Calibrated to the Italian production network, existing disclosure recovers 0.67 percentage points of GDP, predominantly through the markup channel. Aggregate cost stabilization yields little additional welfare because lower cost volatility crowds out disclosure relationships, offsetting most of the direct gain.

Keywords: information frictions, production networks, inflation expectations, supplier disclosure, 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 are linked through prices, quantities, and input demand, so imperfect information about aggregate conditions can distort bilateral allocations and propagate through the production structure, as in La'O and Tahbaz-Salehi (2022). A central friction is that buyers often have to infer suppliers' cost conditions from prices. Suppliers may instead disclose that information directly, removing the need for price-based inference. Whether buyer-supplier relationships operate under price inference or direct disclosure shapes markups, the transmission of shocks through production networks, and, ultimately, welfare. This paper provides causal evidence on these two modes of transmission and quantifies their aggregate implications.

We combine survey, administrative, and experimental data from Italy in what is, to our knowledge, the first setting to link a persistent randomized information treatment to the actual input-output transactions of surveyed firms. This allows us to observe and track supply-chain relationships by their economic importance, duration, and network position over time. Our setting merges the *Survey on Inflation and Growth Expectations* (SIGE), a quarterly Bank of Italy firm survey, with administrative records on 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, together with the transaction values, duration, and network position of each link. 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 become observable to trading partners through repeated interaction.

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 source from broader upstream networks. Thus, direct learning from partners' beliefs and indirect learning from partners' prices move in opposite directions across relationship margins. Downstream learning from customers operates through the belief channel on average, but its heterogeneity across relationship margins is weaker, less stable across specifications, and less clearly tied to a common mechanism than on the supplier side—a pattern the model explains as a consequence of the non-strategic nature of downstream information flow. Crucially, our fixed-effect structure absorbs the common component of partner signals, so that these coefficients are identified from idiosyncratic, partner-specific variation in beliefs and prices.

Our second result is that this heterogeneity is fundamentally sender-side. The gradients are strong and systematic when measured using sender-side relationship characteristics—relationship persistence, economic importance, and supplier breadth (the number of the supplier's upstream input suppliers)—but weaken sharply, and in some cases disappear, when the same characteristics are measured on the receiver side. This asymmetry points to supplier-originated disclosure as the key margin. Whether information is transmitted directly or inferred from prices depends primarily on sender-side relationship characteristics—the supplier's persistence, economic importance, and upstream complexity—rather than on the buyer's side of the relationship.

Our third result is that belief-based transmission has real effects. Using administrative transaction data, we show that the belief channel predicts bilateral trade volumes more robustly

than the price channel. A better-informed supplier's expectations more credibly convey the aggregate component of its costs to buyers, reducing the uncertainty that would otherwise depress 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 disclose 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 disclosure about private information, in the spirit of Crawford and Sobel (1982). Suppliers can credibly disclose information about common cost pressures to buyers, but disclosure must be supported by relationship surplus and sufficiently low setup costs. Each link therefore operates under one of two regimes: a *disclosure regime*, in which buyers receive direct information and suppliers charge the full-information markup, or a *price-inference regime*, in which buyers extract information from prices, perceived demand elasticity falls, and markups exceed the full-information benchmark. Disclosure is more likely in durable, high-surplus relationships and harder to sustain when suppliers source from broader upstream networks, consistent with the cross-link heterogeneity documented in the data. Because downstream information flow is non-strategic—sellers learn from buyers' demand rather than through a disclosure choice—the model predicts that downstream transmission carries no gradient across relationship margins, consistent with the flat downstream heterogeneity documented above. The model also predicts that the belief channel should be the dominant margin for bilateral trade volumes, a prediction we confirm in the data. To our knowledge, this is the first paper to model endogenous bilateral disclosure as an equilibrium margin in a production network, identify that margin empirically, and quantify its aggregate welfare implications.

Calibrated to the Italian production network, the model implies that in the absence of any belief channels the economy incurs a welfare cost of 0.87% of GDP, predominantly through the markup distortion arising from price-based inference. Existing information transmission eliminates three-quarters of this cost, recovering 0.67 percentage points of GDP. The markup distortion accounts for roughly half of the residual equilibrium cost, while the standard second-order pricing errors of La'O and Tahbaz-Salehi (2022) account for the remainder. This markup cost is first order in the informativeness of bilateral prices and responds weakly to aggregate price-level stabilization: the direct effect of removing all aggregate cost variance is a 0.09 pp reduction, but crowding out of disclosure relationships offsets two-thirds of this gain, leaving a net reduction of only 0.03 pp. The welfare response is non-monotone but the differences across stabilization levels are small — the dominant message is that aggregate stabilization has limited traction on the bilateral disclosure problem.

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 Lucas (1972). Fang et al. (2025) study firms' endogenous attention to other firms in an information network, and Huo et al. (2025) show that the structure of such attention networks can amplify micro-level shocks into aggregate fluctuations. These papers take the network of information flows as the object of study, with attention determining which signals firms observe. Afrouzi (2024) studies strategic inattention among horizontal competitors, showing that firms optimally neglect aggregate conditions when competitors do the same. Our paper is complementary: rather than attention allocation, we focus on the endogenous choice between direct disclosure and price-based inference within vertical buyer–supplier relationships.

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. Coibion et al. (2020) provide evidence from Italian firms, showing that treated firms adjust their pricing decisions in response to improved inflation expectations. Weber et al. (2025) extend this to multiple countries and time periods, showing that information treatments become less effective in higher-inflation environments because firms are more attentive and already better informed about inflation. We show that, in our 2019–2023 sample, the treatment nevertheless continues to generate enough exogenous variation in firms' inflation expectations to identify transmission across network links.

Albagli et al. (2022) provide related evidence, using Chilean firm data merged with VAT and customs records, that firms use input prices from their suppliers to update inflation expectations. Our paper is also related to recent work studying how macroeconomic information travels across production networks. Hajdini et al. (2025) use a randomized information treatment in New Zealand to show that beliefs and actions spill over across buyer–supplier pairs and that direct communication is the dominant transmission mechanism on average. Our paper asks a distinct question: whether information is transmitted through suppliers' direct disclosure or through buyers' inference from suppliers' prices, and how the relative importance of these channels varies across relationships. Combining a persistent randomized treatment with the universe of bilateral transactions in Italy, we show that this heterogeneity is driven primarily by supplier-side characteristics and interpret it through a model in which disclosure is an endogenous relationship-level regime.

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) characterizes when market-clearing prices in a network aggregate dispersed private information, showing that network topology governs the speed and completeness of price-based information diffusion, 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 public and private observations of others' actions, more public information can crowd out long-run social learning and reduce asymptotic knowledge. Our contribution is to combine these insights in a framework where the mode of information transmission itself is endogenous. Our evidence shows that both price-based learning and direct disclosure operate in production networks, and our model determines which channel dominates on each link. The price-inference regime that Nikolakoudis (2024) characterizes theoretically is empirically active in our data—particularly on links with suppliers that have broad upstream input networks—and

our framework nests it as the equilibrium outcome on non-disclosure links.

More broadly, our paper relates to work on information frictions, coordination, and shock propagation in macroeconomics, including Angeletos and Huo (2021), and to the literature on production networks and propagation, including Acemoglu et al. (2012), Acemoglu and Tahbaz-Salehi (2025), Gabaix (2014), and Ghassibe and Nakov (2025); see Carvalho and Tahbaz-Salehi (2019) for a survey. On welfare measurement in production networks, Afrouzi et al. (2024) provides a framework for quantifying misallocation costs that arise from information frictions propagating through input-output linkages; our welfare decomposition builds on that tradition while introducing a new first-order term from the price-inference markup distortion. 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 sender-side and receiver-side gradients, and reports an additional state-dependent pattern in information transmission. Section 5 develops the model of strategic information disclosure in production networks and tests its bilateral trade prediction in Section 5.5. Section 6 quantifies the disclosure mechanism: it calibrates the model to the experimental moments, evaluates non-targeted validation checks, 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 allows 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.[1] 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

[1] 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.

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.[2] 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.[3] 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, approximately 14 percent of firms were reassigned from treatment to control; treatment status is otherwise unchanged through 2023.[4] 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 plots average two-quarter-ahead CPI inflation expectations in the SIGE survey for treated and untreated firms together with realized CPI inflation. Treated firms' expectations track realized inflation more closely than those of untreated firms over most of the sample. During the rapid disinflation of 2023–2024, untreated firms appear temporarily closer to realized inflation, largely because they had previously underpredicted the inflation surge and therefore benefited, in levels, from the reversal in inflation. Over the full sample, however, treated firms display persistently smaller forecast errors. This persistent difference in forecast accuracy motivates our analysis of whether the information advantage created by the treatment propagates through production networks.

[2]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.

[3]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.

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

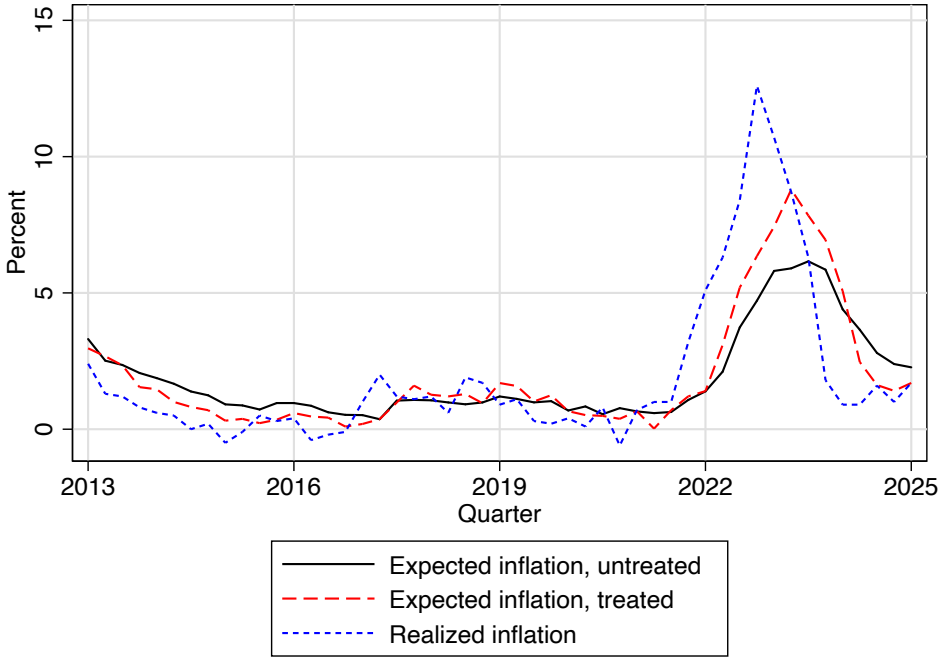


Figure 1: Inflation expectations two quarters ahead and realized inflation

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.[5] 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 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

[5]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.

growth downstream—and allow all spillover coefficients to vary by direction throughout the empirical analysis.

The analysis sample for the expectation-formation regressions 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. Relationship characteristics—duration, economic importance, and supplier breadth—are constructed from the full universe of transactions of each surveyed firm, not only from the matched-pair subsample, exploiting the administrative data's coverage of all bilateral transactions for SIGE participants. 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 matched sample contains 179,464 distinct trading links observed in SIGE, yielding 358,928 directional observations after stacking upstream and downstream directions.

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	1,020	3.00	5.00	9.00	1.00	3.00	6.00	0.2	1.5	11.1
2020	1,223	3.00	6.00	10.00	2.00	3.00	6.00	0.2	1.4	10.0
2021	1,236	4.00	7.00	12.00	2.00	4.00	8.00	0.1	1.0	7.4
2022	1,251	4.00	7.00	11.00	2.00	4.00	8.00	0.2	1.2	7.9
2023	1,242	4.00	6.00	11.00	2.00	5.00	8.00	0.1	1.1	7.8
Downstream										
2019	780	2.00	3.50	8.00	1.00	2.00	5.00	0.1	1.1	7.8
2020	910	2.00	3.00	8.00	1.00	2.00	5.00	0.1	1.1	7.8
2021	961	2.00	4.00	9.00	1.00	2.00	6.00	0.1	1.1	7.8
2022	1,004	2.00	4.00	10.00	1.00	2.00	7.00	0.1	1.1	7.8
2023	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.

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 5–7 upstream suppliers and 3–4 downstream customers, with firms in the upper quartile linked to 9–12 upstream and 8–10 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.

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.

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.[6] 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.

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 selective disclosure mechanism developed in Section 5.

[6]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.

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.[7] 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 ($h=2$).[8] Columns 1–3 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.

[7] 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.

[8] Similar results obtain at alternative horizons.

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 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.[9]

3.2 Learning through beliefs and prices

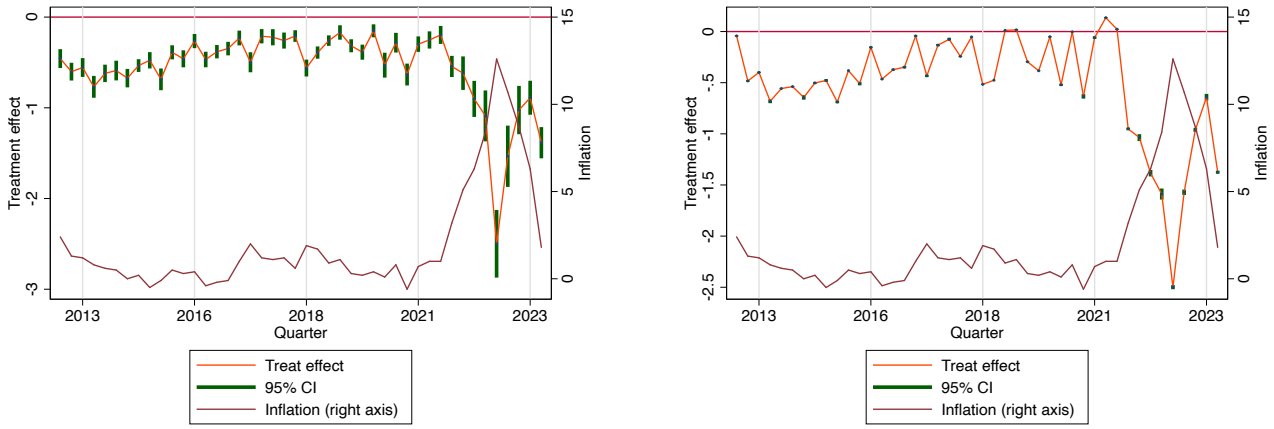
We now test whether firms learn from better-informed trading partners, and through which channel. We distinguish between two modes of transmission. The first is a belief channel, in which firms update their inflation expectations in response to partners' stated beliefs. The second is a price channel, 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. As a preliminary check, treatment reduces the dispersion of inflation

[9]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.

expectations between buyer–supplier pairs, compressing the information gap across linked firms (Appendix B). The analysis in this section refines this finding by separately identifying the belief and price channels through which the information advantage propagates. 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. The positive price-channel coefficients nonetheless confirm that firms do extract information from partners' prices, providing empirical support for theories of price-based information diffusion in networks such as Nikolakoudis (2024). 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.[10]

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

[10]Purchase growth $g^r_{j,t}$ is the year-on-year growth rate of the total taxable amount (*imponibile*) invoiced to customer j, computed from the administrative transaction data.

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.[11] 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 revenues generated 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. Since transaction records contain values rather than quantities and prices separately, $g^r_{j,t}$ reflects both price and quantity changes in the customer's purchases.

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 partners equally. We report robustness to three alternatives. Breadth weights (*supA*; $s_{ij} \propto n^{\text{dir}}_j$, 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-

[11] Appendix L reports horizon-specific estimates.

connected partners, who aggregate information from a larger set of counterparts. Transaction-value weights (*imp*; $s_{ij} \propto \bar{v}_{ij}$, the bilateral taxable amount cumulated over the full sample) upweight economically larger relationships. Relationship-intensity weights (*impJ*; $s_{ij} \propto \bar{v}_{ij}/\bar{V}_j$, the buyer's total purchases from j divided by j's total sales, both cumulated over the full sample) give more weight to links in which firm i accounts for a larger share of j's revenue. All four weighting schemes use full-sample averages and are therefore time-invariant. In all 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. 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 four weighting schemes: uniform (unif), breadth-weighted (supA, based on the supplier's number of upstream input suppliers), transaction-value weighted (imp), and partner-sales-share weighted (impJ, the buyer's share of the supplier's total sales). 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.[12]

[12]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—

Table 3: Learning from supply-chain partners

	Untreated only ($T_{i,t}=0$)				Full sample			
	(1) unif	(2) supA	(3) imp	(4) impJ	(5) unif	(6) supA	(7) imp	(8) impJ
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**	0.040	0.034*	0.024	0.069**	0.030
	(0.014)	(0.020)	(0.029)	(0.030)	(0.017)	(0.017)	(0.027)	(0.027)
δ_2: $T_{i,t}\, \pi^{e,h}_{j,t}$					0.031	0.026	0.038	-0.001
					(0.032)	(0.026)	(0.030)	(0.035)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					-0.034*	-0.027	-0.044*	-0.014
					(0.019)	(0.020)	(0.022)	(0.023)
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***	0.069**	0.056***	0.098***	0.089**	0.065*
	(0.021)	(0.024)	(0.030)	(0.031)	(0.017)	(0.032)	(0.032)	(0.037)
δ_2: $T_{i,t}\, \pi^{e,h}_{j,t}$					0.054**	0.077**	0.071**	0.069*
					(0.019)	(0.033)	(0.031)	(0.037)
δ_3: $T_{i,t}\, T_{j,t}\, \pi^{e,h}_{j,t}$					-0.066**	-0.098**	-0.066	-0.074
					(0.025)	(0.046)	(0.042)	(0.045)
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	0.082

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 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 (0.069 to 0.113, Panel B, columns 1–4) and remains positive in the full-sample specifications (0.056 to 0.098, columns 5–8). On supplier links, δ_1 is also positive and statistically significant under uniform and transaction-value weights in both blocks (0.026 to 0.060 in columns 1–4 and 0.034 to 0.069 in columns 5–8, Panel A), and positive but imprecisely estimated under breadth and relationship-intensity weights. 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. The positive δ_1 is identified from partner j's deviation from the cell mean within each direction–treatment-pair–quarter group. This rules out that the treatment effect on partner beliefs operates only through a common channel—for instance, all treated firms revising in the same direction and pulling their partners along via correlated industry shocks. Firm i loads positively on the *idiosyncratic* component of partner j's belief, indicating genuine extraction of partner-specific information rather than exposure to a common shift.

Price-based inference is weaker and noisier. Learning from transaction-based signals is 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 belief transmission being a cleaner and less noisy signal 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 upstream belief-channel coefficients are stable across weighting schemes, with estimates ranging from 0.026 to 0.060. The downstream belief coefficient shows more variation (0.069–0.113), with larger estimates under transaction-value and breadth-weighted specifications, indicating that downstream transmission is more sensitive to link weighting than the upstream channel. 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.

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 specification estimated separately at $h = 6$ and $h = 12$, and separate-direction estimates at $h = 12$.

Appendix I reports two complementary assessments of magnitude. A variance decomposition shows that partner beliefs add 0.18 percentage points to R^2, concentrated in the customer direction. A linear counterfactual that assigns the treatment to all currently untreated partners implies a mean expectation shift of 0.041 percentage points, roughly 4% of the typical quarterly update. These magnitudes are modest, as expected given that the treatment shifts information precision for a subset of partners and the regression identifies only the idiosyncratic partner-specific component of that shift. Section 6 embeds this mechanism in the full production network to assess its aggregate implications.

These average effects mask substantial heterogeneity across relationships. We turn to that evidence next.

4 Supplier heterogeneity in information transmission

This section examines how the balance between belief-based and price-based learning varies across relationships. 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 time-averaged, 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 $\widetilde{r}_{ij,t}$:

$$\begin{aligned}\pi_{i,t}^{e,h} = {} & \alpha_i + \alpha_h + \alpha_{s(i)\times t} + \alpha_{d(i)\times t} + \alpha_{a(i)\times t} + \alpha_{T_{j,t}\times t} + \rho\, \pi_{i,t-1}^{e,h} + c_1 \widetilde{r}_{ij,t} + c_2 T_{j,t}\, \widetilde{r}_{ij,t} \\ & + \Big[\beta_0 + \beta_1 T_{j,t} + \beta_2 \widetilde{r}_{ij,t} + \beta_3 (T_{j,t} \times \widetilde{r}_{ij,t})\Big]\, g_{j,t} \\ & + \Big[\delta_0 + \delta_1 T_{j,t} + \delta_2 \widetilde{r}_{ij,t} + \delta_3 (T_{j,t} \times \widetilde{r}_{ij,t})\Big]\, \pi_{j,t}^{e,h} + \varepsilon_{ij,t}, \end{aligned} \tag{7}$$

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 $\widetilde{r}_{ij,t}$ is the demeaned log of the heterogeneity measure. In column (1), the partner-side value share $s_{ij}^{(j)} \equiv \bar{v}_{ij} / \sum_{i'} \bar{v}_{i'j}$, where $\bar{v}_{ij}$ is the sum of bilateral taxable amounts across all quarters the dyad is active. In column (2), the partner-side persistence share, defined as the number of active quarters of dyad (i, j) divided by the total number of active quarters across all of j's relationships. In column (3), partner breadth, measured as the time-averaged number of j's upstream input suppliers. 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). Relationship durability, not economic size, 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.[13]

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 suppliers source from broader upstream networks, prices replace beliefs as the main vehicle of information transmission. This pattern is consistent with broad-network suppliers being harder to disclose credibly because their costs depend on many upstream input relationships, 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, the defining feature of downstream learning is the absence of structured heterogeneity across relationship margins, consistent with the model's prediction that downstream information flow is non-strategic.[14]

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.

Sender-side versus receiver-side heterogeneity

The preceding results measure relationship importance from the partner's perspective. Appendix M reports the analogous estimates under receiver-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

[13]Persistence is measured as the dyad's share of partner j's total relationship quarters—a supplier-side measure that captures how central the relationship is to the supplier's network. The buyer-side analog (the dyad's share of firm i's total relationship quarters) is statistically indistinguishable from zero under all specifications, consistent with the supplier-side asymmetry documented in Section 4. The supplier-side gradient is robust to alternative inference: the t-statistic is 5.7 under Driscoll–Kraay, 2.9 under supplier-level clustering, and 2.1 under two-way clustering by buyer and supplier.

[14]The downstream transaction signal captures revenue growth rather than quantity growth separately, since administrative records contain bilateral values rather than prices and quantities independently. This may attenuate estimated coefficients on the transaction-based channel toward zero, but does not affect the belief-based heterogeneity results, which are estimated from the belief coefficient and its gradient across relationship margins.

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 time-averaged 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.

sign, under receiver-side definitions.

This reversal is most striking for persistence. In the joint value-persistence specification, the upstream belief interaction δ_3 is 0.022 (s.e. 0.006, $p < 0.01$) when persistence is measured from the supplier's side, but -0.007 (s.e. 0.010) and insignificant when measured from the receiver's side. The same pattern appears for value-based size: sender-side definitions generate strong or significant upstream belief gradients, whereas receiver-side definitions produce flat, weak, or reversed gradients. For breadth, the sender-side gradient is strong and significant (-0.022, s.e. 0.010), while the receiver-side gradient is positive and only marginally significant (0.033, s.e. 0.018), suggesting that buyer-side network complexity may independently raise the value of supplier belief signals without overturning the dominant sender-side pattern.[15]

This asymmetry is more naturally explained by supplier disclosure incentives than by buyer-side attention or information 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 receiver's perspective. It is equally difficult to reconcile with a pure relationship-quality story: if durable relationships communicate more because of unobserved match quality or trust, the gradient should appear symmetrically on both sides of the relationship, not exclusively on the sender side. 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 suppliers with broad upstream input networks. Second, these gradients are fundamentally sender-side rather than receiver-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] Higher aggregate uncertainty plausibly raises the value of partner-specific information by increasing the dispersion of cost shocks that firms need to forecast; empirically, π_t^2 is a natural proxy be-

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

[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.

cause 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 $\pi^{e,h}_{i,t}$ (percentage points); the full estimating equation is in Appendix R.1. The key coefficients are the belief-channel interactions with uncertainty:

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

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., more than half 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. 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, baseline belief transmission is already stronger, leaving less room for uncertainty to amplify it further. This state dependence is complementary to the finding in Coibion and Gorodnichenko (2025) that firms pay more attention to inflation when the inflation environment is more volatile. Our results point to one concrete source of this heightened attention: firms extract more information from trading partners' beliefs when aggregate uncertainty is high.

[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).

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. Industry-pair equilibrium objects—Kalman weights ω_{ij}, markup distortions $\tilde{\mu}_{ij}$, and residual uncertainties γ_{ij} (all defined in Section 5.2)—depend only on signal precisions and are therefore time-invariant; we suppress the t subscript on these objects throughout. Link-level objects and the disclosure set $\mathcal{R}_t$ vary over time as matches turn over.

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 2). 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 four 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 2). Section 5.3 derives the link-level incentive constraint that governs supplier disclosure (Proposition 3): disclosure is sustainable only on links where the relation-

ship is sufficiently durable and valuable, generating cross-link heterogeneity in upstream belief transmission. Section 5.4 maps the model's channels to the empirical estimates and states the corollary on belief transmission (Corollary 1). Section 5.5 tests the model's bilateral trade prediction: the belief channel, not the price channel, should predict trade volumes conditional on industry controls. Appendix E 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 calibrates the model and quantifies the welfare consequences of imperfect disclosure.

The general equilibrium closure—household preferences, monetary policy, and market clearing—follows La'O and Tahbaz-Salehi (2022) and is in Appendix F. The key simplification is that monetary policy pins down the nominal wage, so individual firms treat it as exogenous; the bilateral results (Propositions 2 and 3) and their empirical predictions are therefore independent of the closure.

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, with productivity decomposed into an industry-common shock $\epsilon_{i,t} \sim \mathcal{N}(0, \sigma^2_{\epsilon,i})$—capturing sector-specific technology or factor-market disruptions shared by all firms in i—and a firm-specific shock $\xi_{k,t} \sim \mathcal{N}(0, \sigma^2_{\xi,i})$, mutually independent across firms and independent of all aggregate variables. We place no restriction on the correlation structure of $\{\epsilon_{i,t}\}_{i=1}^n$ across industries: common shocks are in general correlated because they propagate through the production network. Output is:

$$y_{ik,t} = e^{\epsilon_{i,t}+\xi_{k,t}} \, l_{ik,t}^{\alpha_i} \prod_{j=1}^{n} m_{ijk,t}^{a_{ij}}, \tag{8}$$

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

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. Throughout, we use industry indices (i, j) for objects common to all matches between industries i and j, and firm indices (k, ℓ) for match-specific objects. 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

[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. With finitely many matched suppliers, oligopolistic competition reduces but does not eliminate the markup distortion: each seller still faces a perceived demand elasticity depressed by price-based inference, so the disclosure incentive remains first order. The complication is that

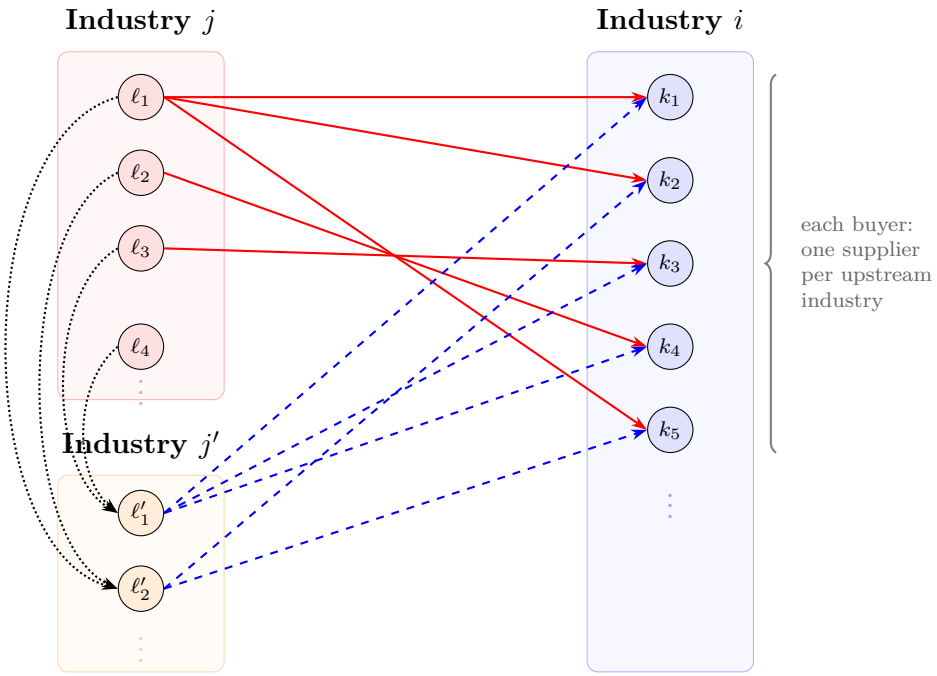


Figure 3: Bilateral matching. Each buyer in industry i draws one supplier per upstream industry (j and j'); draws are independent across buyers. Supplier degree varies: ℓ_1 serves three buyers in i, ℓ_3 one, ℓ_4 none. Red solid arrows show links from industry j to buyers in i; blue dashed arrows show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

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 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 $P_{j,t} \equiv \bar{P}_j \, e^{z_{j,t}}$, where $\bar{P}_j$ is the steady-state wholesale price of industry j and $z_{j,t}$ is the industry-common cost component defined below in (10). 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

suppliers would interact strategically in both pricing and disclosure decisions toward the same buyer, requiring an equilibrium characterization of the joint disclosure set across co-suppliers; we leave this extension for future work.

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.

Given the matching structure, firm k's input costs depend on the prices charged by its matched suppliers $\{\ell_j(k)\}$. The log nominal marginal cost is, to a first-order approximation around the symmetric steady state,[19]

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

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

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

The object $z_{i,t}$ 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 $z_{i,t}$ is a linear combination of jointly Gaussian shocks, the vector $\mathbf{z}_t \equiv (z_{1,t}, \ldots, z_{n,t})'$ is jointly Gaussian with mean zero and covariance matrix $\Sigma_z \equiv \mathrm{Var}(\mathbf{z}_t)$. The diagonal elements $\sigma^2_{z,i} \equiv \mathrm{Var}(z_{i,t})$ are industry-specific; the off-diagonal elements $\mathrm{Cov}(z_{i,t}, z_{j,t})$ 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

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

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]

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

where $\bar{p}_{ij}$ is the steady-state log bilateral price (defined precisely below as $\log \mu_i^* + \overline{mc}_j$), treated as a known constant. The residual $\zeta_{ik,t}$ collects all firm-specific and match-specific deviations of marginal cost from its industry-common component. 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] 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 $z_{i,t} = \sum_j h_{ij}\alpha_j(\log w_t - \bar{w}) + \sum_j h_{ij}\,\epsilon_{j,t}$, where h_{ij} is the (i,j) element of the Leontief inverse $H = (I - A)^{-1}$. The covariance matrix is $\Sigma_z = H\,\Sigma_\epsilon\,H'$, where Σ_ϵ is the covariance matrix of $\{\epsilon_{i,t}\}$. This expression treats the wage w_t as deterministic (or absorbed into the industry fixed effects); with stochastic wages, Σ_z includes an additional term $H\,\alpha\alpha'\,\sigma_w^2\,H'$ where $\alpha = (\alpha_1, \ldots, \alpha_n)'$, which adds a rank-one component to Σ_z proportional to σ_w^2.

[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 $\overline{mc}_j$.

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

(i) Shocks and private signals. Nature draws the vector of common components $\mathbf{z}_t = \{z_{i,t}\}$ and idiosyncratic shocks $\{\xi_{k,t}\}$. Each firm k in industry i perfectly observes its own $\xi_{k,t}$ and receives a noisy private signal about the common component of each industry j:

$$s_{z,k,j,t} = z_{j,t} + \nu_{k,j,t}, \qquad \nu_{k,j,t} \sim \mathcal{N}(0,\, \chi_{ij}\sigma^2_{\nu,j}),$$

independent across firms and industries, where $\chi_{ij} \geq 1$ scales the noise of firm k's signal about industry j relative to the seller's own noise.[22] The noise $\chi_{ij}\sigma^2_{\nu,j}$ is allowed to vary with the industry pair, with $\chi_{ii} = 1$ for own-industry signals; the special case $\chi_{ij} = 1$ for all (i, j) nests the information structure of La'O and Tahbaz-Salehi (2022). The price-flexibility coefficient $\phi_i \equiv \sigma^2_{z,i}/(\sigma^2_{z,i} + \sigma^2_{\nu,i})$ measures how aggressively industry i's firms pass common-component uncertainty into their prices. Each firm forms a prior belief about its marginal cost from these signals.

(ii) Disclosure. Let $\mathcal{R}_t \subseteq \{(k, \ell) : a_{ij} > 0\}$ denote the *disclosure set*, determined endogenously below. Denote by $\hat{z}_{j\ell,t} \equiv \phi_j\, s_{z,\ell,j,t}$ the posterior of supplier ℓ about its own industry common component $z_{j,t}$, where $\phi_j \equiv \sigma^2_{z,j}/(\sigma^2_{z,j} + \sigma^2_{\nu,j})$ is the Kalman weight on its private signal (the analogue of ϕ_i in stage (i) but for the supplier's industry j). On each link $(k, \ell) \in \mathcal{R}_t$, supplier ℓ sends a message $m_{k\ell,t} = \hat{z}_{j\ell,t}$ to buyer k. On non-disclosure links, no voluntary communication occurs.

(iii) Pricing and transactions. We conjecture—and verify in Proposition 1 below—that each firm ℓ in industry j sets its bilateral price as a markup over its expected marginal cost. This log-linear pricing rule is exact up to Jensen's inequality corrections that are second order in shock variances:

$$\log p_{k\ell,t} \;=\; \log \mu_{ij} \;+\; \hat{z}_{j\ell,t} \;+\; \hat{\zeta}_{j\ell,t} \;+\; \overline{mc}_j, \tag{13}$$

where $\hat{\zeta}_{j\ell,t} \equiv \mathbf{E}_{j\ell}[\zeta_{j\ell,t}]$ is the seller's prior about the idiosyncratic component of its marginal cost, and μ_{ij} is the equilibrium markup derived in Section 5.2.[23] Prices are set using beginning-of-period information only: own $\xi_{\ell,t}$ and signal $s_{z,\ell,j,t}$, before upstream input prices are observed.[24]

Buyer k's demand for supplier ℓ's good is

$$x_{kj,t}(p) \;=\; \left(\frac{p}{\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]}\right)^{-\varepsilon_i} A_{kj,t}, \tag{14}$$

where $A_{kj,t}$ is a demand shifter, $\varepsilon_i > 1$ is the demand elasticity, and $\mathcal{I}_{k,t}$ is the buyer's information set at the time of the purchase decision. On *disclosure links*, $\mathcal{I}_{k,t}$ contains the supplier's disclosed belief $\hat{z}_{j\ell,t}$; on *non-disclosure links*, it contains the bilateral price $p_{k\ell,t}$, from which the buyer extracts information about $z_{j,t}$ via Bayesian updating. The price-based inference reduces the buyer's perceived demand elasticity and generates the markup distortion derived in Section 5.2.[25]

[22] Firm k perfectly observes $\xi_{k,t}$ for simplicity. The model's key mechanism survives without it: the firm-specific residual $\zeta_{ik,t}$ arises endogenously from match-specific price variation regardless of $\xi_{k,t}$.

[23] The seller's expectation of its idiosyncratic cost component equals the firm-specific shock minus a correction that depends on its upstream price signals: $\hat{\zeta}_{j\ell,t} = \xi_{\ell,t} - \sum_r a_{jr}\, \phi_r(1 - \phi_r)\, s_{z,\ell,r,t}$.

[24] Following La'O and Tahbaz-Salehi (2022), we treat the seller's signal as exogenous to the pricing game, abstracting from models such as Nikolakoudis (2024) where prices aggregate information across network layers.

[25] The demand shifter is $A_{kj,t} \equiv a_{ij} S_{i,t}/\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$, where a_{ij} is the input-output share and $S_{i,t}$ is the buyer's total nominal expenditure on inputs, both known to k. 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). The full derivation is in Appendix D.

Supplier ℓ observes the quantity demanded, which provides a second signal about $z_{j,t}$ (the buyer's signal $s_{z,k,j,t}$ is independent of the seller's), but arrives after pricing and is payoff-irrelevant under i.i.d. shocks.[26]

(iv) End-of-period updating. Each firm aggregates all signals from the period into a posterior belief.[27]

5.2 Bilateral pricing and surplus

We solve the bilateral problem by backward induction through the pricing and belief formation stages 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 E and Appendix F, respectively.

Price-based inference. On non-disclosure links, the buyer uses the supplier's price to update beliefs about $z_{j,t}$. By the pricing rule (13), the price residual

$$\log p_{k\ell,t} - \bar{p}_{ij} - (\log \mu_{ij} - \log \mu_i^*) \;=\; \hat{z}_{j\ell,t} + \hat{\zeta}_{j\ell,t} \tag{15}$$

is a noisy signal about $z_{j,t}$: it confounds the common component with the idiosyncratic term $\hat{\zeta}_{j\ell,t}$, just as in Lorenzoni (2009).[28][29] The buyer's posterior mean therefore satisfies

$$\log \mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}] \;\approx\; \log \bar{P}_j \;+\; \omega_{ij}\Big(\log p_{k\ell,t} - \bar{p}_{ij} - (\log \mu_{ij} - \log \mu_i^*)\Big), \tag{16}$$

so a one-unit increase in $\log p$ raises $\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$ by ω_{ij}, partially offsetting the demand reduction. The supplier internalizes this when choosing its markup. The Kalman weight is:[30]

$$\omega_{ij} = \frac{\gamma_{ij}}{\gamma_{ij} + \sigma^2_{\zeta,j}}, \qquad \gamma_{ij} \equiv \mathrm{Var}(\hat{z}_{j\ell,t} \mid s_{z,k,j,t}), \tag{17}$$

where $\sigma^2_{\zeta,j}$ is the variance of the supplier's idiosyncratic component $\hat{\zeta}_{j\ell,t}$ (the noise in the price residual, analogous to $\sigma^2_{\zeta,i}$ in (11) but for the supplier's industry j), and γ_{ij} is the buyer's residual uncertainty about the seller's posterior after using its private signals. Under $\chi_{ij} \to \infty$ for $j \neq i$, the buyer has no useful independent information about industry j prior to observing the price, so γ_{ij} depends on the buyer's own-industry signal $s_{z,k,i,t}$ only through the cross-industry covariance $\mathrm{Cov}(z_j, z_i)$, and ω_{ij} is common across all buyers in industry i facing suppliers in j—hence the (i,j) subscript.[31]

[26] Online Appendix H introduces attention signals that add a baseline downstream information channel insensitive to relationship characteristics, matching the empirical patterns in Section 4. The core bilateral mechanism does not depend on them.

[27] Under i.i.d. $z_{i,t}$, this posterior is payoff-irrelevant: next period's shock is independent. With persistent shocks, the posterior becomes the prior for $t+1$ and bilateral learning affects future pricing. The i.i.d. case isolates the within-period demand channel from the dynamic pricing channel.

[28] The key extension is that disclosure eliminates this confounding on links where the supplier communicates $\hat{z}_{j\ell,t}$ directly; the model determines when this is credible.

[29] Under independent matching, $\hat{\zeta}_{j\ell,t}$ is orthogonal to the buyer's information set, so ω_{ij} in (17) fully summarizes the informativeness of bilateral prices.

[30] The equilibrium markup μ_{ij} depends on ω_{ij} (derived below), determined by signal precisions and the disclosure configuration $\mathcal{R}_t$, not on shock realizations.

[31] A Kalman calculation gives $\gamma_{ij} = \phi_j^2(\sigma^2_{z,j} + \sigma^2_{\nu,j}) - \phi_j^2 \mathrm{Cov}(z_j, z_i)^2/(\sigma^2_{z,i} + \sigma^2_{\nu,i})$, which depends on the IO network through $\mathrm{Cov}(z_j, z_i)$. On disclosure links, $m_{k\ell,t} = \hat{z}_{j\ell,t}$ is sufficient, so $\mathrm{Var}(\hat{z}_{j\ell,t} \mid s_{z,k,i,t}, m_{k\ell,t}) = 0$ and

Pricing without inference. On disclosure links, $\mathcal{I}_{k,t}$ contains $\hat{z}_{j\ell,t}$ directly, so demand (14) has constant elasticity ε_i and the price carries no additional information. Define the markup $\mu \equiv p/mc_{j\ell,t}$, the demand scale $\Gamma_{k\ell,t} \equiv mc_{j\ell,t}^{1-\varepsilon_i} A_{kj,t}$, and the per-unit profit function $\pi(\mu) \equiv (\mu-1)\mu^{-\varepsilon_i}$, so $\Pi^S_{k\ell,t}(\mu) = \pi(\mu)\,\Gamma_{k\ell,t}$. The profit-maximizing markup satisfies $\pi'(\mu_i^*) = 0$, which gives

$$\mu_i^* = \frac{\varepsilon_i}{\varepsilon_i - 1}, \tag{18}$$

with $\pi''(\mu_i^*) < 0$, confirming a maximum. The pricing rule (13) reduces to $\log p^*_{k\ell,t} = \bar{p}_{ij} + \hat{z}_{j\ell,t} + \hat{\zeta}_{j\ell,t}$, where $\bar{p}_{ij} \equiv \log \mu_i^* + \overline{mc}_j$ (with $\overline{mc}_j$ the supplier's steady-state log marginal cost from (11)), and $\pi^*_{ij} \equiv \pi(\mu_i^*)$ denotes the frictionless per-unit profit.[32]

Pricing with inference. Since a fraction ω_{ij} of each price increase is absorbed by the belief shift rather than by lost demand, the effective demand elasticity falls from ε_i to $\varepsilon_i(1-\omega_{ij})$. The supplier's problem on non-disclosure links is therefore

$$\max_{\mu} \; (\mu - 1)\,\mu^{-\varepsilon_i(1-\omega_{ij})}\,\widetilde{\Gamma}_{k\ell,t}, \tag{19}$$

where $\widetilde{\Gamma}_{k\ell,t} \equiv mc_{j\ell,t}^{\varepsilon_i\,\omega_{ij}}\,\Gamma_{k\ell,t}$ is the demand scale inflated by the factor $mc_{j\ell,t}^{\varepsilon_i\,\omega_{ij}} > 1$: price-based inference effectively raises the supplier's perceived demand by making buyers less responsive to price increases. The first-order condition yields the equilibrium markup

$$\tilde{\mu}_{ij} = \frac{\varepsilon_i(1-\omega_{ij})}{\varepsilon_i(1-\omega_{ij}) - 1} > \mu_i^*, \tag{20}$$

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

$$\tilde{\mu}_{ij} - \mu_i^* = \frac{\varepsilon_i\,\omega_{ij}}{(\varepsilon_i - 1)^2} + O\left(\omega_{ij}^2\right). \tag{21}$$

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

The profit gain from non-disclosure for the seller is

$$\Delta\tilde{\pi}_{ij,t} \equiv \tilde{\pi}_{ij} - \pi^*_{ij} = \pi^*_{ij} \cdot \varepsilon_i \log \mu_i^* \cdot \omega_{ij} + O\left(\omega_{ij}^2\right), \tag{22}$$

where $\tilde{\pi}_{ij} \equiv (\tilde{\mu}_{ij} - 1)(\tilde{\mu}_{ij})^{-\varepsilon_i(1-\omega_{ij})}$ is the seller's profit at the distorted markup and elasticity. The gain is *first order* in ω_{ij}: 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.

Proposition 1 (Equilibrium pricing)**.** *The equilibrium bilateral log price takes the form* (13)*, where:*

(i) *On non-disclosure links:* $\mu_{ij} = \tilde{\mu}_{ij} = \varepsilon_i(1-\omega_{ij})/[\varepsilon_i(1-\omega_{ij}) - 1] > \mu_i^*$.

(ii) *On disclosure links:* $\omega_{ij} = 0$ *and* $\mu_{ij} = \mu_i^* = \varepsilon_i/(\varepsilon_i - 1)$.

Proof. Parts (i) and (ii) follow from the first-order conditions (18) and (20). See Lemma 1 in Appendix C for the full derivation including Jensen's inequality corrections. □

$\omega_{ij} = 0$. Under a regularity condition on the IO structure (Online Appendix G), ω_{ij} is increasing in supplier breadth, consistent with Section 4.

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

5.3 Disclosure

The markup distortion on non-disclosure links generates a profit gain for the seller that voluntary information transmission can eliminate. We now ask: in which direction does disclosure generate a seller profit gain large enough to justify the cost of sustaining the information channel?

Proposition 2 (Asymmetric value of directional information)**.** *Consider a bilateral link between supplier ℓ in industry j and buyer k in industry i, and let $\omega_{ij} \in (0,1)$ denote the informational content of the bilateral price (equation* (17)*).*

(i) *Upstream (seller → buyer). The seller's profit gain from upstream disclosure is first order in ω_{ij}: $\tilde{\pi}_{ij} - \pi^*_{ij} = \pi^*_{ij} \cdot \varepsilon_i \log \mu^*_i \cdot \omega_{ij} + O(\omega^2_{ij})$, 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 $O(\omega^2_{ij})$—second order in the same perturbation parameter.*

Proof. See Appendix C. □

The economic intuition follows from the envelope condition. The buyer's surplus is $\Pi^B_{k\ell,t}(\mu) = \mu^{1-\varepsilon_i}\Gamma_{k\ell,t}/(\varepsilon_i - 1)$, so total surplus $W_{k\ell,t} \equiv \Pi^S_{k\ell,t} + \Pi^B_{k\ell,t}$ satisfies[33]

$$W'_{k\ell,t}(\mu) = \varepsilon_i \, \mu^{-\varepsilon_i - 1}(1-\mu)\, \Gamma_{k\ell,t}, \tag{23}$$

so $W'_{k\ell,t}(\mu^*_i) < 0$: any markup above marginal cost destroys total surplus. And since $\pi'(\mu^*_i) = 0$ by (18), a small markup change has only a second-order effect on the seller's profit but a first-order effect on buyer surplus. On a non-disclosure link the seller charges $\tilde{\mu}_{ij} > \mu^*_i$, exploiting the buyer's lower perceived elasticity $\varepsilon_i(1-\omega_{ij})$. Disclosing eliminates this exploitation: the seller loses $\tilde{\pi}_{ij} - \pi^*_{ij}$ today, which is first order in ω_{ij} because $\tilde{\mu}_{ij}$ is optimal at a different elasticity than μ^*_i. For the downstream direction, the seller already charges its privately optimal markup μ^*_i where $\pi'(\mu^*_i) = 0$; improving the seller's information about $z_{j,t}$ 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.[34][35]

Importantly, upstream disclosure also generates a *first-order welfare gain*: moving from $\tilde{\mu}_{ij}$ to μ^*_i raises total surplus $W_{k\ell,t}$ by an amount proportional to $\tilde{\mu}_{ij} - \mu^*_i = O(\omega_{ij})$, because $W'_{k\ell,t}(\mu^*_i) < 0$ by (23). The gain falls on the buyer: since $\pi'(\mu^*_i) = 0$, the seller's profit loss equals

[33] Online Appendix G derives Π^B from the CES demand integral. Away from the steady state, $\Gamma_{k\ell,t}$ depends on beliefs; welfare comparisons are evaluated to first order in δ, holding $\Gamma_{k\ell,t}$ at its steady-state value.

[34] The asymmetry requires only $\pi'(\mu^*_i) = 0$ (private markup optimum) and $\tilde{\mu}_{ij} \neq \mu^*_i$ (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.

[35] The seller's incentive to overstate its belief—which raises the buyer's perceived outside option and insulates demand from price increases—rules out informative cheap-talk equilibria in the upstream direction by the logic of Crawford and Sobel (1982), motivating relational enforcement as the mechanism that sustains disclosure. In the downstream direction, the buyer would need to disclose its demand conditions to the seller. But the buyer always prefers a lower input price and therefore has a uniform incentive to understate demand, ruling out informative cheap talk by the same Crawford–Sobel logic. The second-order profit gain from downstream disclosure makes relational enforcement unprofitable regardless, so the downstream channel remains non-strategic.

the deviation gain $\tilde{\pi}_{ij} - \pi^*_{ij}$, while the buyer's surplus gain from the lower markup is strictly larger, so total surplus rises.

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 E, these attention signals provide the baseline downstream information channel on all links.

Relational enforcement. Although upstream disclosure generates a first-order profit gain for the seller, it is not self-enforcing: a seller who has been disclosing can profitably deviate by stopping disclosure this period, charging the distorted markup $\tilde{\mu}_{ij}$, and earning $\tilde{\pi}_{ij} > \pi^*_{ij}$ today. Sustaining disclosure therefore requires a commitment technology. We model this as a relational contract: the buyer detects deviations and permanently reverts to price-based inference upon detection, so the seller faces a permanent future loss of $\pi(\mu^*_i) - \pi(\tilde{\mu}_{ij}) > 0$ per period, where $\pi(\mu) \equiv (\mu - 1)\mu^{-\varepsilon_i}$ is the per-unit profit at the frictionless elasticity. Establishing this channel requires a one-time setup cost $\psi \geq 0$, capturing the investment needed to make cost reports credible—agreeing on a communication protocol, building the mutual familiarity that makes stated beliefs verifiable, or setting up informal accounting transparency.

IC condition. Disclosure is verifiable: the buyer detects non-disclosure immediately and reverts to price-based inference in the same period. The seller therefore cannot exploit a one-period surprise gain. Disclosure is established if and only if the present value of disclosing forever outweighs both the setup cost and the present value of reverting forever:

$$\psi \;\leq\; R_{ij}\left[\pi(\mu^*_i) - \pi(\tilde{\mu}_{ij})\right]\Gamma_{k\ell,t}, \tag{24}$$

where $R_{ij} \equiv 1/(1 - \beta\rho_{ij})$ is the capitalization factor (including the current period), and $\pi(\mu^*_i) - \pi(\tilde{\mu}_{ij}) > 0$ because μ^*_i maximizes $\pi(\mu)$ at the frictionless elasticity ε_i. The right-hand side is the PV of the per-period profit gain from operating at the frictionless markup μ^*_i rather than the distorted markup $\tilde{\mu}_{ij}$ that would prevail under non-disclosure.[36]

The IC condition (24) characterizes which links sustain disclosure as a function of model primitives. Because the right-hand side rises with ρ_{ij} (through R_{ij}) and with $\sigma^2_{z,j}$ (through ω_{ij}), disclosure is a threshold phenomenon in relationship durability: each link discloses if and only if ρ_{ij} is high enough to make the future surplus outweigh the setup cost, with the threshold falling as industry uncertainty rises.

Figure 4 illustrates the IC condition. The green curve plots the present value of the per-period profit gain from disclosure, $R[\pi(\mu^*) - \pi(\tilde{\mu}(\omega))]$, where $\pi(\mu)$ is evaluated at the frictionless elasticity. The red horizontal line is the setup cost $\psi/\bar{\Gamma}$. The green curve rises with ω because a higher price distortion means a larger gain from switching to the frictionless markup. The threshold ω^* is where the two curves cross; links to the right disclose.

Proposition 3 (Selective disclosure)**.** *For each link (k, ℓ) with $k \in i$ and $\ell \in j$, there exists a unique threshold $\bar{\rho}_{ij} \in (0, 1)$ such that disclosure is established if and only if $\rho_{ij} \geq \bar{\rho}_{ij}$. The threshold is:*

[36] Grim-trigger is the most severe credible punishment and therefore maximizes the set of sustainable disclosure relationships; any less severe punishment would only tighten the IC condition. The buyer's reversion to price-based inference is credible because it is individually rational. See Appendix C for the full derivation.

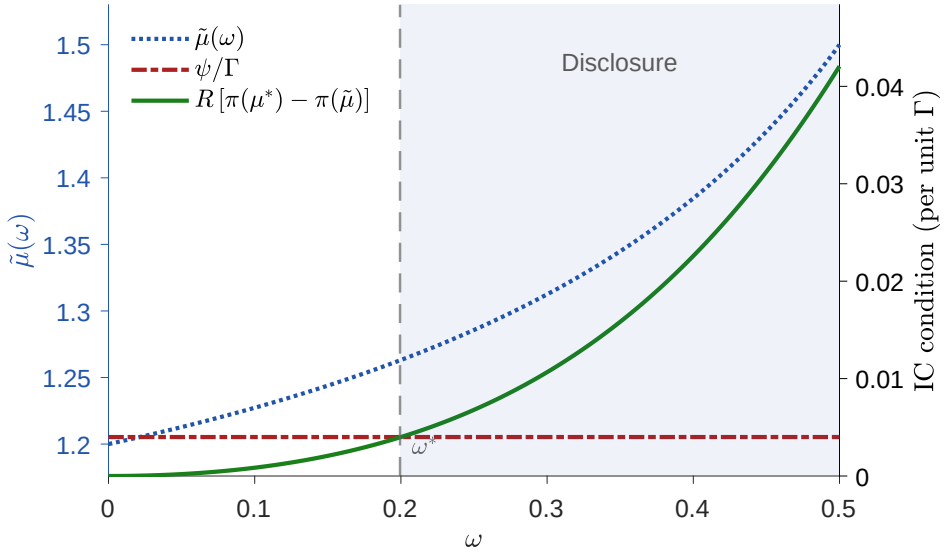


Figure 4: The disclosure IC condition

Notes: Left axis (blue): distorted markup $\tilde{\mu}(\omega) = \varepsilon(1-\omega)/[\varepsilon(1-\omega)-1]$. Right axis: the two sides of the IC condition (24) per unit of bilateral trade $\bar{\Gamma}$. Red: setup cost $\psi/\bar{\Gamma}$ (horizontal line). Green: present value (PV) of future profit loss from reversion, $R\,[\pi(\mu^*) - \pi(\tilde{\mu}(\omega))]$, where $\pi(\mu) \equiv (\mu-1)\mu^{-\varepsilon}$ is evaluated at the frictionless elasticity ε and $\pi^* \equiv \pi(\mu^*) > \pi(\tilde{\mu}(\omega))$ since μ^* maximizes $\pi(\mu)$. The threshold ω^* is where the two curves cross: links with $\omega > \omega^*$ satisfy the IC condition and disclose (shaded region). $\varepsilon = 6$; R and $\psi/\bar{\Gamma}$ are illustrative.

(a) *decreasing in $\sigma^2_{z,j}$ (aggregate cost volatility): higher industry-level uncertainty makes the price signal more informative, raising the gains from disclosure;*

(b) *increasing in $\sigma^2_{\zeta,j}/\sigma^2_{z,j}$ (idiosyncratic-to-aggregate noise ratio): noisier bilateral prices reduce the Kalman weight ω_{ij} and hence the markup distortion, lowering the gains from disclosure;*

(c) *decreasing in bilateral trade volume $\Gamma_{k\ell,t}$: larger trade stakes raise the PV of future gains from charging μ_i^*.*

The proof is in Appendix C.

Finally, the mechanism delivers the state dependence central to the empirics. Higher aggregate cost volatility—whether from wage volatility σ_w^2 or industry-specific shocks $\sigma^2_{\epsilon,j}$—raises $\sigma^2_{z,j}$ through the input-output network (10), lowers $\bar{\rho}_{ij}$, and activates disclosure on links previously below threshold. This is consistent with the activation of the belief channel during the 2021–2023 inflation episode documented in Section 4, when a surge in input costs raised cost uncertainty across the production network.

5.4 Mapping to the empirical evidence

The model delivers a precise mapping between structural objects and the regression coefficients estimated in Sections 3–4. The mapping depends on the timing of the survey relative to trade realization: if the expectation is elicited before all trades within a period have occurred, it reflects the posterior conditional on the subset of links that have already transacted. The corollary below remains valid because it is stated at the link level—each realized link contributes its

marginal weight to the posterior (characterized below), and the qualitative ordering across links is preserved regardless of which subset has realized. The empirical estimates in Table 3 identify these weights by exploiting cross-link variation in treatment status, averaging over the realized subsets in the data.

Corollary 1 (Belief transmission channels)**.** *Consider a bilateral link (k, ℓ) with $k \in i$ and $\ell \in j$. The marginal contribution of supplier ℓ's information to buyer k's posterior—holding all other signals fixed—operates through three channels:*

(i) Upstream, disclosure links *($\rho_{ij} \geq \bar{\rho}_{ij}$): loads on $\hat{z}_{j\ell,t}$ directly, with weight ϕ_j (the fraction of $z_{j,t}$ reflected in j's belief).*

(ii) Upstream, non-disclosure links *($\rho_{ij} < \bar{\rho}_{ij}$): loads on the price residual* (15) *with weight $\omega_{ij}\phi_j$.*

(iii) Downstream*: supplier ℓ learns from buyer k's information set $\mathcal{I}_{k,t}$. The weight on this channel is non-strategic and relationship-invariant.*

In the single-link case these coincide with $\kappa_{ij}(1) = \phi_j$ and $\kappa_{ij}(0) = \omega_{ij}\phi_j$; in the multi-link network the weights become partial regression coefficients $\kappa_{ij}(T_j)$ that also depend on cross-signal correlations from other links.

Proof. Immediate from (27), Proposition 3, and (14). □

The weights $\kappa_{ij}(T_j)$ are partial regression coefficients from the joint projection of $z_{i,t}$ on the full signal vector, and therefore depend on the precision of every signal in the economy—not only on the bilateral ω_{ij}. The treatment differential $\kappa_{ij}(1) - \kappa_{ij}(0)$ identifies the belief channel in Table 3, with the bilateral pricing problem determining how ω_{ij} feeds into $\kappa_{ij}(0)$ (Appendix A). The upstream gradient varies with relationship characteristics (Proposition 3); the downstream channel is non-strategic and relationship-invariant, consistent with the flat gradient in Section 4, Panel B.[37]

5.5 A test of the model: bilateral trade

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 ($T_i = 0$):

$$\begin{aligned}\log \Gamma_{k\ell,t} = \alpha_k + \delta_{s(j),t} + \alpha_{s(k),t} + \alpha_{\text{size}(k),t} + \alpha_{\text{area}(k),t} + \alpha_{T_j,t} \\ + \beta_0^p\, g_{j,t}^p + \beta_1^p\, (T_j \times g_{j,t}^p) + \beta_0^e\, \pi_{j,t}^e + \beta_1^e\, (T_j \times \pi_{j,t}^e) + \varepsilon_{k\ell,t},\end{aligned} \tag{25}$$

[37] The downstream belief coefficient in Table 3 is consistent with item (iii): a treated buyer's reported expectation is a natural proxy for its information set $\mathcal{I}_{k,t}$, the object the supplier learns from, so the belief interaction captures the downstream channel directly.

where $\Gamma_{k\ell,t}$ is the transaction value between buyer k and supplier ℓ in quarter t, $g^p_{j,t}$ is the supplier's price growth, and $\pi^e_{j,t}$ 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 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 model delivers sharp predictions from the CES demand equation (14). Demand depends on the buyer's expected price index $\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$, which in turn depends on the information available at the time of purchase. On disclosure links, the buyer's information set contains the supplier's communicated belief $\hat{z}_{j\ell,t}$; on non-disclosure links, it contains only the bilateral price. Treatment improves the informativeness of both signals, generating $\beta^p_1 > 0$ (price channel) and $\beta^e_1 > 0$ (belief channel). The model predicts an additional asymmetry: conditional on supplier-industry×quarter fixed effects, which absorb the common component of industry prices, the price channel should be attenuated while the belief channel survives. This is because the price signal operates through industry-level variation in ω_{ij}, which is absorbed by industry fixed effects, whereas the disclosed belief varies at the supplier level and is not absorbed. The dominant prediction is therefore $\beta^e_1 \gg \beta^p_1$ in specifications with industry controls (Appendix R.2).

Results. Table 7 reports the results. The belief channel is robust and economically dominant across all weighting schemes, while the price channel is weak and not robust once industry controls are included. The treated-supplier belief interaction $\hat{\beta}^e_1$ is large and highly significant across all four specifications, while the treated-supplier price interaction $\hat{\beta}^p_1$ ranges from small and insignificant to marginally significant depending on the weighting scheme, with no robust pattern.

Under uniform weights (column 1), $\hat{\beta}^e_1 = 0.104$ (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 ($\hat{\beta}^e_0 \approx -0.10$): 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 ($\hat{\beta}^e_0 + \hat{\beta}^e_1 \approx 0$).

This pattern has a natural interpretation within the model, and maps directly onto the model's central mechanism. The trade evidence is primarily about persistent relationship-level differences in informativeness: links with better-informed suppliers sustain systematically higher trade volumes because treatment status and disclosure regimes are persistent, not because of quarter-to-quarter variation in survey responses. Higher supplier inflation expectations normally

signal higher expected costs, which would 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.

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 trade result should therefore be read primarily as evidence on cross-link differences in disclosure regimes: links with better-informed suppliers sustain systematically higher trade volumes, consistent with disclosure being a persistent relationship-level margin rather than a within-link dynamic response.

The weakness of the price channel in the bilateral trade data—despite its marginal significance in some specifications and in the expectations regressions—is informative. In the model, the price channel operates through the buyer's Bayesian updating of beliefs about $z_{j,t}$ 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.

Table 7: 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6 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. Finally, we quantify the welfare cost of information frictions, evaluate counterfactual information regimes, and assess the scope for monetary policy to reduce the dominant welfare cost through aggregate price stabilization.

The calibration proceeds in two steps. First, industry-level signal parameters ($\hat{\sigma}^2_{z,j}$, $\hat{\sigma}^2_{\zeta,j}$, $\hat{\sigma}^2_{\nu,j}$) and link-level characteristics (ρ_{ij}, Γ_{ij}) are measured directly from the data. Second, two scalar parameters (ψ and λ_ζ) are matched to two aggregate moments in an exactly-identified GMM system, with all remaining empirical patterns serving as non-targeted validation checks (Table 9, Panel B).

6.1 Calibration

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

External parameters. Industry-level signal parameters are measured directly from the survey and price data. Table 8 summarizes their cross-industry distributions.

The variance of the common component is $\hat{\sigma}^2_{z,j} = \mathrm{Var}_t(\bar{g}^p_{j,t})$, 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:

$$\bar{g}^p_{j,t} \;=\; \mu_j \;+\; \phi_j\,\bar{c}_{j,t} \;+\; \eta_{j,t}, \tag{26}$$

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 $\hat{\phi}_j$ estimates the Kalman gain $\phi_j = \sigma_z^2/(\sigma_z^2+\sigma_\nu^2)$, from which we recover $\hat{\sigma}^2_{\nu,j} = \hat{\sigma}^2_{z,j}(1-\hat{\phi}_j)/\hat{\phi}_j$. 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 z_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 ϕ_j (and hence σ_ν^2) from the correct information set.[38]

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.

At median industry values, the price–cost pass-through is $\hat{\phi} = 0.48$, implying that sellers pass roughly half of cost movements into prices. The implied private-signal noise is $\hat{\sigma}_\nu = 2.44$ pp at the median, substantially larger than the common-component standard deviation ($\hat{\sigma}_z = 2.32$ pp), confirming that the private signal is noisy. The aggregate price pass-through is $\theta'\phi = 0.50$, 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. At the median, the raw ratio $\hat{\sigma}_\zeta/\hat{\sigma}_z \approx 1.9$ implies that idiosyncratic price variation substantially exceeds the common component in the data. This ratio will be rescaled by λ_ζ in the internal calibration below; the raw ratio provides a lower bound on the signal-to-noise problem facing buyers.

Disclosure costs. The model uses a uniform setup cost ψ across all links. The calibration implements the IC condition (24) with two adjustments for aggregation. First, the match-specific trade volume $\Gamma_{k\ell,t}$ is replaced by the industry-pair average $A_{ij} \equiv \mathbf{E}[\Gamma_{k\ell,t}]$, since $\Gamma_{k\ell,t}$ is

[38] As a testable implication of own-signal pricing, $\hat{\phi}_j$ should not differ by treatment status, since treatment affects the buyer's precision but not the seller's signal. We find $\hat{\phi}^{\text{ut}} = 0.52$ and $\hat{\phi}^{\text{tr}} = 0.49$, with an interaction of -0.03 (s.e. 0.04), indistinguishable from zero.

Table 8: External parameters

Parameter	Symbol	Median	P10	P90
Panel A: Signal standard deviations (pp, 78 industries)				
Common-component std. dev.	$\hat{\sigma}_{z,j}$	2.32	1.10	5.39
Idiosyncratic price noise	$\hat{\sigma}_{\zeta,j}$	4.48	2.59	6.57
Private-signal noise	$\hat{\sigma}_{\nu,j}$	2.44	1.38	3.21
Panel B: Price–cost pass-through (61 industries)				
Pass-through coefficient	$\hat{\phi}_j$	0.48	0.20	0.71
Pass-through (pooled, ind. FE)	$\hat{\phi}$	0.49 (s.e. 0.03)		
Treatment interaction	$\hat{\phi}^{\mathrm{tr}} - \hat{\phi}^{\mathrm{ut}}$	−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 standard deviations. $\hat{\sigma}_{z,j}$ and $\hat{\sigma}_{\zeta,j}$ are measured from price data; $\hat{\sigma}_{\nu,j}$ is derived from the pass-through coefficient $\hat{\phi}_j$ via $\hat{\sigma}_\nu^2 = \hat{\sigma}_z^2(1-\hat{\phi})/\hat{\phi}$. Panel B reports the price–cost pass-through from regression (26). 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: $\varepsilon = 6$, $\beta = 0.99$.

not observed at the industry-pair level. Second, the per-period surplus $\pi(\mu^*) - \pi(\tilde{\mu}_{ij})$ is evaluated at the distorted markup $\tilde{\mu}_{ij}$ that would prevail on a non-disclosure link, because the IC condition measures the gain from switching into disclosure.

Internally calibrated parameters. The mapping from parameters to moments is transparent. A higher ψ raises the IC threshold for disclosure, reducing the disclosure rate and hence the upstream belief premium (the coefficient on $T_j \times \pi^e_{j,t}$ in Table 3); ψ is therefore identified by this moment. A higher λ_ζ makes prices noisier, reducing the price-signal weight $\omega_{p,j}$ and hence the upstream price coefficient (the coefficient on $T_j \times g^p_{j,t}$ in the same table); λ_ζ is identified by this moment. Figure 5 illustrates both identification relationships.

We set $\chi_{ij} = \infty$ for all $j \neq i$: firms have no useful independent signal about other industries' cost conditions prior to observing prices and disclosed beliefs. This is the natural assumption in a supply chain setting where firms observe their own production costs precisely but learn about upstream and downstream conditions only through trading partners. It implies that ω_{ij} is larger than under partial cross-industry observability, amplifying the markup distortion on non-disclosure links. The welfare numbers are nonetheless disciplined by the data: λ_ζ is calibrated to match the upstream price coefficient, which directly pins down the empirical magnitude of ω_{ij}, keeping the welfare decomposition within data-consistent bounds.

Table 9 reports the fit.

Calibration fit. Both targeted moments are treatment effects on survey inflation expectations, which reflect the end-of-period posterior that incorporates bilateral learning (stage (iv) of the timing protocol). The model counterparts are treatment-induced shifts in the weight the buyer places on the supplier's price signal ($\Delta\omega_{ij}$, the price channel) and on the supplier's disclosed belief ($\Delta\kappa_{ij}$, the belief channel), as characterized in Corollary 1. The regressions measure

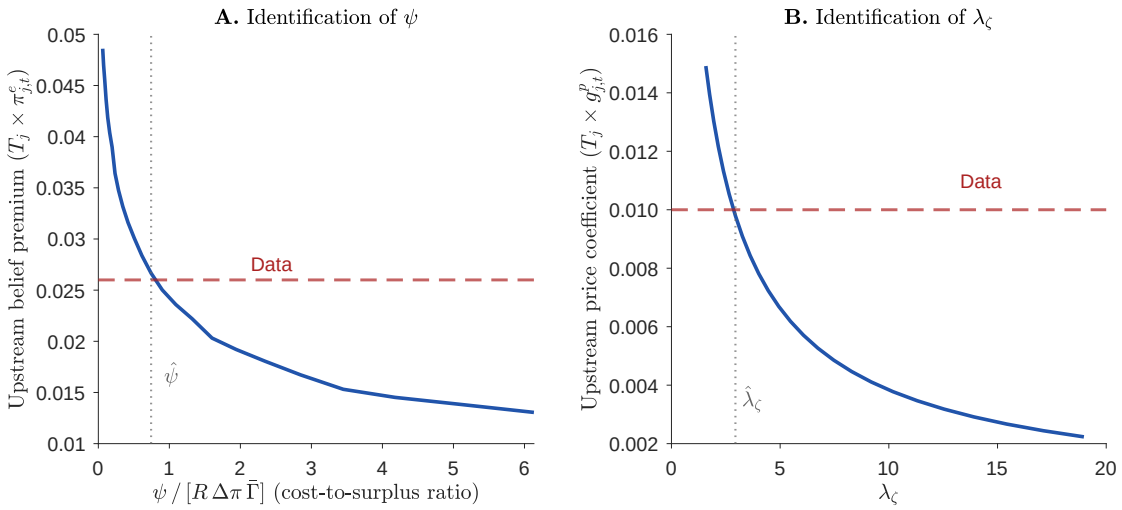


Figure 5: Identification of calibrated parameters

Notes: Panel A plots the model-implied upstream belief premium $(T_j \times \pi^e_{j,t})$ as a function of the cost-to-surplus ratio $\psi \,/\, [R\,\Delta\pi\,\bar{\Gamma}]$, which equals one at the disclosure threshold. Panel B plots the upstream price coefficient $(T_j \times g^p_{j,t})$ as a function of the price-noise scaling λ_ζ. Blue curves are model-implied moments; red dashed lines are the data targets from Table 9. Dotted vertical lines mark the calibrated values $(\hat{\psi},\, \hat{\lambda}_\zeta)$. Each parameter is identified by the monotone relationship between it and its targeted moment, holding the other parameter fixed.

how treated-partner signals enter the buyer's full posterior, and the model computes exactly this object through the bilateral Frisch–Waugh–Lovell decomposition. These moments discipline the demand-side belief system (ψ 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 (26). 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.

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	0.028	0.014
(2) Upstream price coeff. $(T_j \times g^p_{j,t})$	0.010	0.009	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	0.050	0.028
(6) Trade belief coeff. $(T_j \times \pi^e_{j,t}$, trade, lagged)	0.063	0.033	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 5.5. Calibrated parameters: $\lambda_\zeta = 3.16$ ($\hat{\sigma}_\zeta$ scaled from 4.48 to 7.96 pp); ψ normalised to GDP units. External parameters (medians): $\hat{\sigma}_\nu = 2.44$ (from pass-through $\hat{\phi} = 0.48$), $\hat{\sigma}_\zeta = 4.48$, $\hat{\sigma}_z = 2.32$. Fixed: $\varepsilon = 6$, $\beta = 0.99$.

The calibrated ψ is less than one percent of the PDV of bilateral profits $R_{ij}\,\pi^*\,\Gamma_{ij}$ at the median disclosing link: the cooperative surplus from disclosure is itself a small fraction of total bilateral profits, which is why most links remain in the non-disclosure region despite the low

absolute level of the setup cost.

$\lambda_\zeta = 3.16$ implies that the variance of idiosyncratic price noise used in the model is 3.16× the raw within-industry cross-firm variance of price growth measured from the administrative data ($\hat{\sigma}^2_{\zeta,\text{data}}$). This scaling is economically plausible: the within-industry cross-firm variance of price growth averages over many buyer–supplier matches, understating the noise in any single bilateral price signal. A specific buyer observing a single supplier's price faces match-specific pricing variation, relationship-specific factors, and the idiosyncratic component of that supplier's costs, all of which make the bilateral price a noisier signal than the industry average would suggest.At this calibrated value, $\hat{\sigma}_\zeta/\hat{\sigma}_z \approx 1.9$ at the median, so idiosyncratic price variation is nearly twice the common component. This implies a small Kalman weight ω_{ij}: bilateral prices carry limited information about the supplier's common cost component, which explains the small 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. It fits level and size moments well but underfits the persistence gradient, and should therefore be read as conservative on welfare: the true markup distortion from non-disclosure is likely larger than what the calibrated model implies, since a richer persistence mechanism would push more links toward the disclosure threshold. Table 9 reports the full picture.

The bilateral trade belief coefficient (moment 6) is the cleanest out-of-sample test: documented in Section 5.5, it uses a different outcome variable, a different sample, and a different identification strategy from the targeted moments. The model predicts 0.033 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 50% (0.050 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 45% (0.036 vs. 0.079): in the model, downstream transmission is passive demand monitoring (Section 5.3), so the premium reflects treatment-induced precision gain weighted by the Kalman gain on the demand signal.

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 $R_{ij} = 1/(1-\beta\rho_{ij})$, 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 ψ 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.

Equilibrium disclosure. The calibrated model implies that disclosure is substantial but far from universal. Table 10 shows that approximately 26% of active links sustain disclosure, with a strong gradient by link size: 5% in the bottom quartile of expenditure share versus 57% in the top quartile. Disclosure also responds substantially to aggregate uncertainty, rising from 8% to 43% when σ_z^2 quadruples.[39] 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	25.6%
Bottom quartile (expenditure share)	5%
Top quartile (expenditure share)	57%
Low uncertainty ($0.5 \times \bar{\sigma}_z^2$)	8%
High uncertainty ($2 \times \bar{\sigma}_z^2$)	43%

Figure 6 reports the equilibrium disclosure rate by quartile of expenditure share a_{ij}. Disclosure rises steeply with link size: 5% in the bottom quartile versus 57% 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. Proposition 3 predicts that the disclosure threshold $\bar{\rho}_{ij}$ is decreasing in ρ_{ij}, so holding all else equal disclosure is monotone increasing in ρ_{ij}. In the cross-section of calibrated links, however, the unconditional gradient is non-monotone: disclosure rises from near zero at low ρ to above 40% around the median, then declines in the right tail. This reflects a negative correlation between ρ_{ij} and $\Gamma_{k\ell,t}$ in the data: very persistent links tend to have below-average expenditure shares, so the IC condition—which depends on the product $R_{ij} \cdot W(\omega_j) \cdot \Gamma_{ij}$—is not monotone in ρ_{ij} alone. Conditioning on expenditure share restores monotonicity, consistent with Proposition 3.

The CDF in Panel A illustrates the selective nature of disclosure: links that choose to disclose 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.

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

We measure welfare as the consumption-equivalent loss relative to the flexible-price equilibrium, expressed as a fraction of GDP using Domar weights. Following La'O and Tahbaz-Salehi (2022), the welfare cost decomposes into three channels: a *markup channel* ($\mathcal{W}_1$) from distorted

[39] This counterfactual compares $0.5 \times \bar{\sigma}_z^2$ (low uncertainty) to $2 \times \bar{\sigma}_z^2$ (high uncertainty), holding all other parameters—$\sigma_{\zeta,j}^2$, $\sigma_{\nu,j}^2$, ρ_{ij}, Γ_{ij}, ψ—fixed at their calibrated values and re-solving the IC fixed point.

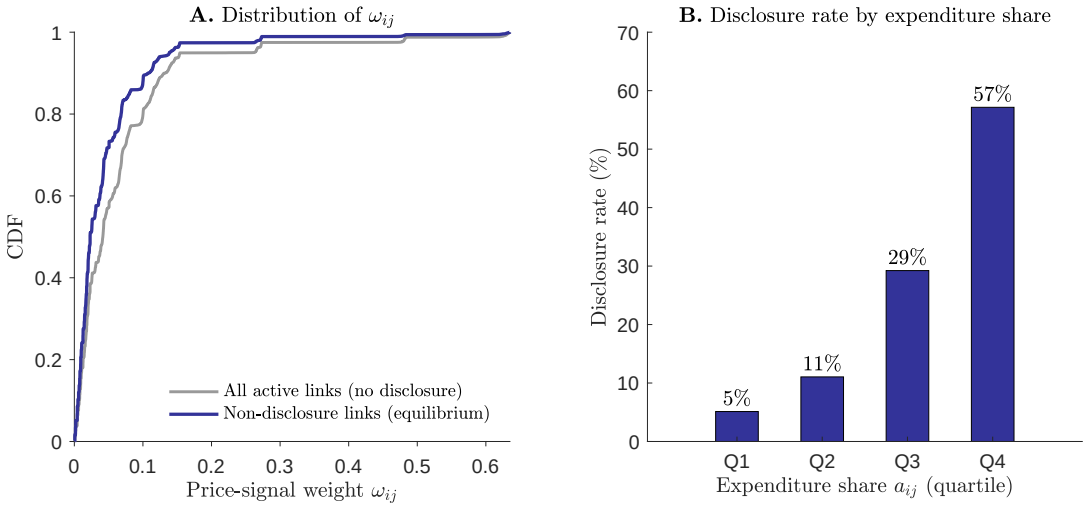


Figure 6: Equilibrium disclosure set

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 4,470 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 equilibrium disclosure rate by quartile of expenditure share a_{ij}. Calibrated parameters as in Table 9.

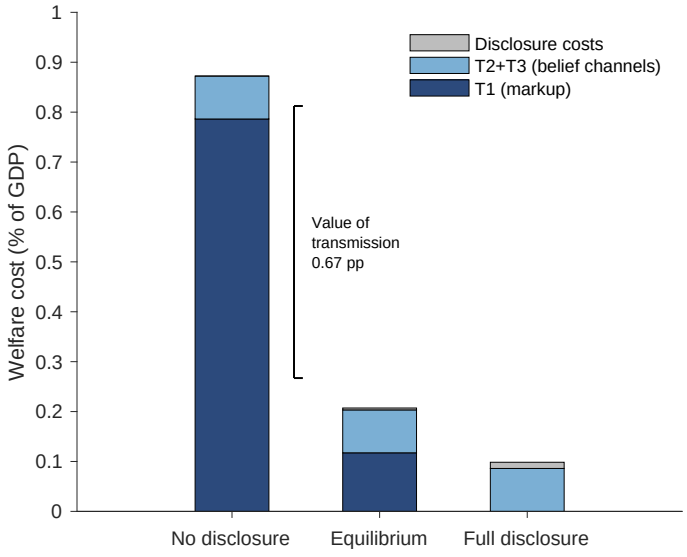


Figure 7: Welfare cost of information frictions

Notes: Consumption-equivalent welfare cost of information frictions decomposed into the first-order markup channel ($\mathcal{W}_1$, dark) and the second-order information channels ($\mathcal{W}_2 + \mathcal{W}_3$, light), across three regimes. The bracket indicates the value of existing information transmission—the welfare gain from moving from the no-belief-channels regime to equilibrium disclosure.

markups on non-disclosure links, an *aggregate belief-error channel* ($\mathcal{W}_2$) from output gap volatility and across-industry misallocation, and a *within-industry misallocation channel* ($\mathcal{W}_3$) from idiosyncratic pricing errors. The markup channel is first order in ω_{ij} and is the novel term introduced by the price-inference mechanism; the remaining two channels are second order in belief errors and correspond to the welfare costs in La'O and Tahbaz-Salehi (2022). The full derivation is in Appendix F.

Welfare decomposition. Figure 7 and Table 11 report the consumption-equivalent welfare cost of information frictions across three information regimes—no belief channels, equilibrium disclosure, and full disclosure. Following La'O and Tahbaz-Salehi (2022), we decompose welfare losses into within-industry misallocation ($\mathcal{W}_3$), output gap volatility and across-industry misallocation ($\mathcal{W}_2$), and a novel first-order markup channel ($\mathcal{W}_1$) arising from price-based inference on non-disclosure links and absent in their framework.[40]

The equilibrium cost is 0.21% of GDP. The majority of this (0.12 pp) arises from the markup channel: non-disclosure links sustain distorted markups that generate first-order welfare losses. Output gap volatility and across-industry misallocation ($\mathcal{W}_2$) contribute 0.026% and within-industry misallocation ($\mathcal{W}_3$) 0.060%—together roughly one-third of the total. Crucially, the output gap and dispersion channels 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 the markup channel and disclosure costs vary with the extent of disclosure. Under the SIGE treatment, the output gap and dispersion channels change slightly because treatment reduces σ_ν^2 for treated firms, which raises ϕ_j and alters the own-signal pricing error; this is distinct from a disclosure regime change. The markup cost is moderately sensitive to the right tail of the ω distribution: capping the markup distortion at $\omega \leq 0.50$ reduces the markup cost by roughly 4%, confirming that the welfare estimate is not driven by extreme links.

Removing all belief channels raises the total cost to 0.87% 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 0.67 percentage points of GDP, overwhelmingly through the markup channel: disclosure eliminates the markup distortion on the 26% of links that sustain it, reducing the markup cost from 0.79% to 0.12%.

The output gap and dispersion channels sum to 0.09% of GDP and are unchanged across regimes, confirming that own-signal pricing insulates them from the disclosure set. The markup cost ranges from 0.79% (no disclosure) to 0.00% (full disclosure)—an order of magnitude larger and the source of the entire value of transmission.

Treatment counterfactual and the markup channel. The SIGE treatment is implemented at the supplier level: when a buyer is matched with a treated supplier, the buyer's private signal noise falls from $\sigma_{\nu,j}^2$ to $0.77 \times \sigma_{\nu,j}^2$ (the treated-to-untreated within-cell forecast dispersion ratio, from Section 6.1). This captures the treatment's effect on the buyer's demand-side precision without altering the pricing equation. The resulting welfare improvement is small (-0.009 pp, i.e., slightly lower costs), reflecting the interaction between the intensive and extensive margins of the price channel.

Treatment improves firms' private precision (ϕ_j rises), which has two opposing effects. On the *extensive margin*, better precision raises the cooperative surplus from disclosure, potentially

[40] A Monte Carlo verification with 100,000 replications confirms that the analytical welfare formula matches simulated welfare to within 5% accuracy; the residual reflects higher-order terms in the Taylor expansion.

Table 11: Welfare cost of information frictions (% of GDP)

	$\mathcal{W}_1$ (markup)	$\mathcal{W}_2$ (agg. error)	$\mathcal{W}_3$ (dispersion)	Disc. costs	Total
Panel A: Information regimes					
No belief channels	0.79	0.03	0.06	0.00	0.87
Equilibrium disclosure	0.12	0.03	0.06	0.00	0.21
Full disclosure	0.00	0.03	0.06	0.01	0.10
Value of transmission	0.67	0.00	0.00	−0.01	0.67
Panel B: Treatment counterfactual					
SIGE treatment	0.12	0.02	0.06	0.00	0.20
Treatment effect	+0.00	−0.01	−0.01	+0.00	−0.01
Links switched			176 / 6,011		
Panel C: Aggregate stabilization counterfactual					
Industry-specific only ($\mathrm{Var}(\bar{z}) = 0$)	0.12	0.02	0.04	0.00	0.18
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 $R = \mathbf{0}$ and $\sigma^2_{\mathrm{dn}} \to \infty$ (suppressing both the disclosed-belief and downstream channels); "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. Note that the markup cost rises slightly under treatment because the intensive margin (higher ω_{ij} on non-disclosure links) partly offsets the extensive margin (links switching into disclosure); see the discussion in Section 6.3. Panel C decomposes each industry's cost shock z_j into a consumption-weighted aggregate component $\bar{z} = \sum_i \theta_i z_i$ and an industry-specific residual. "Industry-specific only" re-solves the model with $\mathrm{Var}(\bar{z}) = 0$, retaining only residual industry-level uncertainty.

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: 176 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 output gap and dispersion channels improve slightly under treatment (-0.008 and -0.005 pp respectively), roughly offsetting the markup cost 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.

6.4 Aggregate stabilization and the limits of monetary policy

The markup distortion on non-disclosure links arises from buyer uncertainty about the common component of the supplier's costs, $z_{j,t}$, which reflects both an aggregate component shared across industries—driven by the nominal wage and correlated productivity shocks—and an industry-specific residual. Monetary policy, by controlling nominal aggregate demand, can in principle stabilize the aggregate component. How much of the equilibrium welfare cost could be eliminated by such aggregate stabilization?

Panel C of Table 11 reports the result. We decompose each industry's common component z_j into a consumption-weighted aggregate $\bar{z} = \sum_i \theta_i z_i$ and an orthogonal industry-specific residual, then re-solve the model with $\text{Var}(\bar{z}) = 0$. Perfect aggregate stabilization reduces the equilibrium welfare cost from 0.21% to 0.18% of GDP—a modest reduction of 0.03 pp, accounting for 13% 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 26% to 14%. 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.

Figure 8 shows how the welfare cost varies continuously with the degree of aggregate stabilization $\alpha \in [0, 1]$, where α is the fraction of $\text{Var}(\bar{z})$ removed. The direct effect (fixing the disclosure set at its baseline equilibrium) is monotonically decreasing in α, as expected. The net effect under endogenous disclosure is non-monotone, reaching a minimum around $\alpha \approx 0.8$, but welfare differences across stabilization levels are small once crowding out is accounted for — the equilibrium welfare cost changes by less than 0.003 pp between $\alpha = 0.8$ and $\alpha = 1$. The dominant message is that the two curves diverge sharply: each additional unit of aggregate stabilization shrinks the bilateral surplus from disclosure, pulling more links below the IC

threshold and reconstituting markup distortions on the newly non-disclosing links. Panel B shows the equilibrium disclosure rate, which declines monotonically with α and accounts for the amplification.[41]

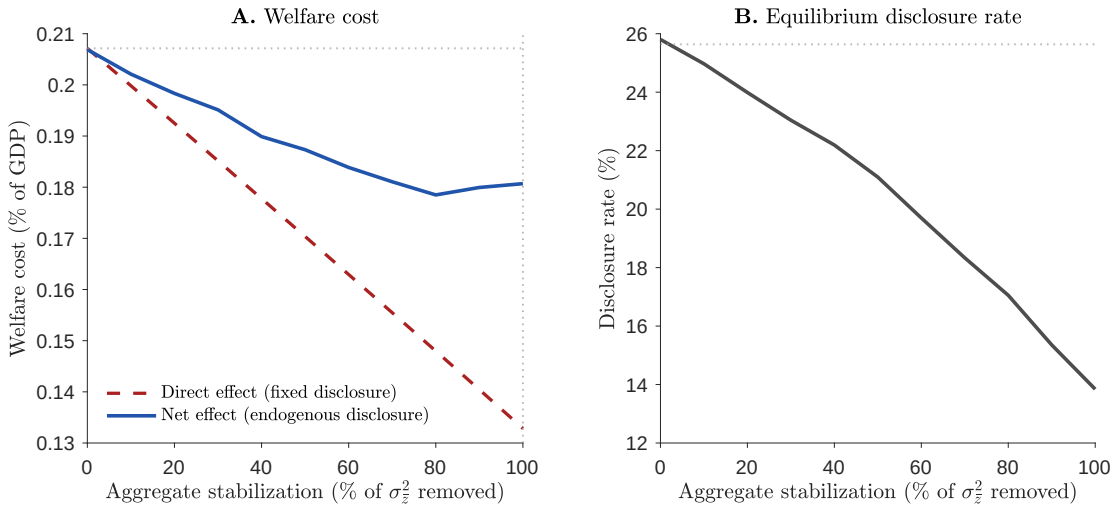


Figure 8: Aggregate stabilization and endogenous disclosure.

Notes: Panel A plots the equilibrium welfare cost (solid blue, endogenous disclosure) and the direct-effect counterfactual (dashed red, fixing the disclosure set at its baseline equilibrium) as a function of α, the fraction of $\text{Var}(\bar{z})$ removed by aggregate stabilization. The horizontal dotted line marks the baseline welfare cost ($\alpha = 0$). Panel B plots the equilibrium disclosure rate. The gap between the two curves in Panel A reflects crowding out of disclosure relationships: as aggregate uncertainty falls, fewer links find disclosure privately optimal, partially offsetting the direct welfare gain.

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 0.87% no-disclosure cost (47%). Without disclosure, there is no margin to crowd out: prices are the only information channel, and removing the aggregate component of σ_z^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.[42]

[41] We verify robustness to equilibrium selection by solving the fixed-point in both directions—sweeping α from 1 to 0 warm-started from the baseline equilibrium, and from 0 to 1 cold-started from no disclosure. The two sweeps yield identical welfare values at every α, confirming that the equilibrium is unique and the pattern is not a warm-start artifact.

[42] This crowding-out 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 differs in two respects. First, the offset operates through the first-order markup channel, rather than through second-order pricing errors. 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, so the crowding-out is socially inefficient.

The implication is that aggregate price stabilization has limited traction on the bilateral communication problem that drives the dominant welfare cost in this economy. The key friction is not the level of aggregate uncertainty per se, but whether trading partners communicate directly or rely on price inference—a margin that monetary policy cannot directly address.

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 that source from broader upstream networks. These gradients are fundamentally sender-side rather than receiver-side, pointing to the supplier's incentives to disclose 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 source from broader upstream networks, consistent with 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 0.67 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 0.21% 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. The breadth gradient documented in the data also points to a richer theory of communication costs: suppliers with more complex upstream cost structures find it harder to credibly disclose to buyers, and formalizing this friction—for instance through a model of verifiable cost reports or disclosure

precision—could generate sharper predictions about which network positions are most exposed to price-inference distortions.

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.

Afrouzi, H., S. Bhattarai, and E. Wu (2024). The welfare cost of inflation in production networks. Mimeo, Columbia University.

Albagli, E., F. Grigoli, and E. Luttini (2022). Inflation expectations and the supply chain. Mimeo, Central Bank of Chile.

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.

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.

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

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.

Appendix

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 (43), firm k in industry i forms a posterior about the common component $z_{i,t}$ as a weighted combination of all available signals:

$$\hat{z}_{ik,t} \;=\; \kappa_i \, s_{z,k,i,t} \;+\; \sum_{j \in \mathcal{N}_i^{\text{up}}} w_{ij,t} \, \tilde{s}_{ij,k,t} \;+\; (\text{downstream and attention signals}), \tag{27}$$

where $\tilde{s}_{ij,k,t}$ denotes the effective signal from upstream partner j: on disclosure links, $\tilde{s}_{ij,k,t} = \hat{z}_{j,t}$ with weight $w_{ij,t} = \psi_{ij,t}$; on non-disclosure links, $\tilde{s}_{ij,k,t} = \hat{z}_{j,t} + \hat{\zeta}_{j\ell_j(k),t}$ (the price residual) with weight $w_{ij,t} = \beta^p_{ij,t}$. All weights are partial regression coefficients from the joint projection of $z_{i,t}$ 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 $\hat{z}_{j,t}$ more precise. This raises the informational value of j's signals from firm i's perspective: because j's posterior now contains more information about $z_{j,t}$ than i can reconstruct from its own signals, the residual uncertainty $\gamma_{ij,t} = \text{Var}(\hat{z}_{j,t} \mid s_{z,k,i,t},\, m_{k\ell,t})$ increases, and the optimal Kalman weight $\omega_{ij,t}$ (equation 17) and belief weight $\beta^p_{ij,t}$ in the full posterior both rise. A second, general-equilibrium channel is that the improved precision of j's signals shifts how j's prices and beliefs covary with other signals in i's information set, altering the entire weight vector.

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

$$\kappa_{ij}(T_{j,t}) \;\approx\; \kappa_{ij}(0) \;+\; T_{j,t}\Big[\kappa_{ij}(1) - \kappa_{ij}(0)\Big]. \tag{28}$$

Substituting into (27), the contribution of partner j's signal to firm i's posterior becomes

$$\kappa_{ij}(0)\, \tilde{s}_{ij,k,t} \;+\; T_{j,t}\Big[\kappa_{ij}(1) - \kappa_{ij}(0)\Big]\, \tilde{s}_{ij,k,t}. \tag{29}$$

The first term is a level loading on the partner's signal; the second is the treatment interaction, which identifies the weight shift induced by the treatment.

Observable signals and the sufficiency of beliefs. The econometrician does not observe the latent signals $\hat{z}_{j,t}$ 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 $\hat{z}_{j,t}$, 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 $z_{j,t}$ and $\zeta_{j\ell,t}$).

A key implication of the model is that $\hat{z}_{j,t}$ is a sufficient statistic for the information about $z_{j,t}$ contained in the price residual: on disclosure links, $\beta^p_{ij,t} = 0$ and the buyer ignores the price signal entirely (equation 43). 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 $z_{j,t}$. 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 idiosyncratic cost shocks $\hat{\zeta}_{j\ell,t}$ and mechanical cost pass-through.

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 $z_{j,t}$ from partner j's stated expectation. The $\boldsymbol{\beta}$ coefficients on $g_{j,t}$, conditional on $\pi^{e,h}_{j,t}$, identify the price channel—the effect of partner price changes on firm i's expectations through idiosyncratic cost shocks and pass-through, net of any information about $z_{j,t}$ already captured by stated 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 κ_{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:

$$f(x_{j,t},\, T_{i,t},\, T_{j,t};\, \boldsymbol{\theta}) \;=\; \theta_0\, x_{j,t} + \theta_1\, T_{j,t}\, x_{j,t} + \theta_2\, T_{i,t}\, x_{j,t} + \theta_3\, T_{i,t}\, T_{j,t}\, x_{j,t}. \tag{30}$$

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 $\hat{z}_{ik,t}$ 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 $\kappa_{ij}(1) - \kappa_{ij}(0)$ from (28). By sufficiency, this is the channel through which learning about $z_{j,t}$ operates.
- $\boldsymbol{\beta}$ coefficients (price channel, $x_{j,t} = g_{j,t}$): conditional on $\pi^{e,h}_{j,t}$, these capture the residual effect of partner price changes on firm i's expectations—idiosyncratic cost shocks and pass-through—through channels other than learning about $z_{j,t}$.
- θ_2 captures the shift in how a treated firm i weights partner signals due to its own improved precision; θ_3 captures the complementarity between own and partner precision—the additional adjustment when both sides of the link are treated.

Role of fixed effects. The direction×treatment-pair×quarter fixed effects $\alpha_{c(i,j)\times t}$ absorb three objects: (i) the direct effect of own treatment on firm i's expectations; (ii) the average effect of partner treatment status on beliefs within each direction–treatment cell; and (iii) aggregate price adjustments that affect all firms in a cell symmetrically when treatment shifts the disclosure set $\mathcal{R}_t$ or industry-level signal precision. 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 partner-specific belief 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

$$y_{ij,t}^{\text{sym}} \equiv \frac{2\left|\pi_{i,t}^{e,h} - \pi_{j,t}^{e,h}\right|}{\left|\pi_{i,t}^{e,h}\right| + \left|\pi_{j,t}^{e,h}\right|}. \tag{31}$$

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

Specification. We estimate

$$y_{ij,t}^{\text{sym}} = \alpha_t + \beta_{TU}\,\mathbf{1}\{TU_{ij,t}\} + \beta_{UT}\,\mathbf{1}\{UT_{ij,t}\} + \beta_{TT}\,\mathbf{1}\{TT_{ij,t}\} + X_{ij,t}'\varpi + \varepsilon_{ij,t}, \tag{32}$$

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

Lemma 1 (Verification of the pricing conjecture)**.** *The log-linear pricing rule* (13) *is consistent with profit maximization. Jensen's inequality corrections do not affect the equilibrium markup $\tilde{\mu}_{ij}$, but they do affect the level of bilateral surplus and hence the value of disclosure; this is a second-order correction to the IC condition* (24) *that vanishes as shock variances shrink.*

Proof. The seller's log nominal marginal cost is $\log mc_{j\ell,t} = \overline{mc}_j + z_{j,t} + \zeta_{j\ell,t}$, where $\overline{mc}_j$ is the steady-state log marginal cost, $z_{j,t}$ is the common component, and $\zeta_{j\ell,t}$ is the idiosyncratic component. The seller sets $\log p = \log \mu + \log mc_{j\ell,t}$ and, using the guessed form (13), replaces $z_{j,t}$ and $\zeta_{j\ell,t}$ with its posteriors $\hat{z}_{j\ell,t}$ and $\hat{\zeta}_{j\ell,t}$:

$$\log p_{k\ell,t} = \log \mu_{ij} + \overline{mc}_j + \hat{z}_{j\ell,t} + \hat{\zeta}_{j\ell,t}.$$

The buyer observes p and extracts a signal about $z_{j,t}$ via (15). Given the Gaussian–linear structure, the buyer's conditional expectation is

$$\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}] = \bar{P}_j \, e^{\hat{z}^B_{k,t} + \frac{1}{2}\operatorname{Var}(z_{j,t}|\mathcal{I}_{k,t})},$$

where $\hat{z}^B_{k,t}$ is the buyer's posterior mean and the variance term is the Jensen correction. The correction depends only on signal precisions, not on shock realizations, so it enters $\widetilde{\Gamma}_{k\ell,t}$ as a deterministic multiplicative constant and drops out of the seller's first-order condition with respect to μ.

The seller's profit-maximization problem on a non-disclosure link is

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

where ω_{ij} is the Kalman weight (17). The first-order condition yields $\tilde{\mu}_{ij}$ as in (20), independently of the Jensen correction. The optimal log price is therefore

$$\log p^*_{k\ell,t} = \log \tilde{\mu}_{ij} + \overline{mc}_j + \hat{z}_{j\ell,t} + \hat{\zeta}_{j\ell,t},$$

which is exactly the conjectured form (13) with $\mu_{ij} = \tilde{\mu}_{ij}$. On disclosure links, $\omega_{ij} = 0$ and $\mu_{ij} = \mu^*_i$, again matching the conjecture. The fixed point is therefore consistent, confirming the equilibrium. □

Proof of Proposition 2. **Part (i).** Expand $\tilde{\pi}_{ij,t} = (\tilde{\mu}_{ij,t} - 1)(\tilde{\mu}_{ij,t})^{-\varepsilon_i(1-\omega_{ij,t})}$ around $\omega_{ij,t} = 0$. At $\omega = 0$: $\tilde{\mu} = \mu_i^*$ and $\tilde{\pi}_{ij} = \pi_{ij}^*$. Differentiating with respect to ω at $\omega = 0$:

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

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

$$\left.\frac{\partial \tilde{\pi}}{\partial \omega}\right|_{\omega=0} = (\mu^*-1)(\mu^*)^{-\varepsilon} \cdot \varepsilon \log \mu^* = \pi_{ij}^* \cdot \varepsilon_i \log \mu_i^* \;>\; 0.$$

Hence $\tilde{\pi}_{ij,t} - \pi_{ij}^* = \pi_{ij}^* \cdot \varepsilon_i \log \mu_i^* \cdot \omega_{ij,t} + O(\omega^2)$, 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 $\hat{z}_j$ is $\hat{z}_j$ 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 $z_{j,t}$). The per-period expected profit is

$$\mathbf{E}\Big[\Pi^S_{k\ell,t}\Big(p^*_{k\ell}(\hat{z}_j)\Big)\Big] = \pi_{ij}^*\,\Gamma_{k\ell,t} - \tfrac{1}{2}|\pi''(\mu_i^*)|\;\mathrm{Var}(z_{j,t} \mid \mathcal{I}_{j\ell})\,\Gamma_{k\ell,t} + O\Big(\mathrm{Var}(z_{j,t} \mid \mathcal{I}_{j\ell})^2\Big),$$

where the second term captures the expected loss from pricing error (the log price deviates from the frictionless target by $(z_{j,t} - \hat{z}_j)$, whose variance is $\mathrm{Var}(z_{j,t} \mid \mathcal{I}_{j\ell})$). 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

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

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

$$\Delta\sigma^2_{\text{post}} = \frac{1}{\iota_j} - \frac{1}{\iota_j + \kappa_{\text{add}}} = \frac{\kappa_{\text{add}}}{\iota_j(\iota_j + \kappa_{\text{add}})}, \tag{33}$$

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

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

which is bounded above by $|\pi''(\mu_i^*)|\Gamma_{k\ell,t}/(2\iota_j^2)$ 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

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

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.

□

Proof of Proposition 3. Define $R_{ij} \equiv 1/(1-\beta\rho_{ij})$ and write the establishment condition (24) as a function of $\sigma_z^2 \equiv \sigma_{z,j}^2$, holding all other link characteristics fixed.

Step 1: Scaling of the cooperative surplus. Since $W'_{k\ell,t}(\mu) = \varepsilon\mu^{-\varepsilon-1}(1-\mu)\,\Gamma_{k\ell,t}$ by (23), we have from (22):

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

For small σ_z^2, the informational content of the price satisfies $\omega_{ij,t} \approx \gamma_{ij,t}/\sigma_{\zeta,j}^2$ (from (17) for small $\gamma_{ij,t}$). Since $\gamma_{ij,t} \le \mathrm{Var}(\hat{z}_{j\ell,t}) = \sigma_{z,j}^4/(\sigma_{z,j}^2+\sigma_{\nu,j}^2) \le \sigma_{z,j}^2$, we have $\omega_{ij,t} = O(\sigma_z^2/\sigma_{\zeta,j}^2)$. Hence

$$\Delta W_{k\ell,t} = \underbrace{\frac{\varepsilon\,|W'_{k\ell,t}(\mu_i^*)|/\Gamma_{k\ell,t}}{(\varepsilon-1)^2\sigma_{\zeta,j}^2}}_{\equiv\ c_W > 0}\ \sigma_z^2\,\Gamma_{k\ell,t}\ +\ O\big(\sigma_z^4\big)\,. \tag{34}$$

Step 2: Threshold from the disclosure IC. The IC (24) requires

$$\psi\ \le\ \Pi_{ij,t}^{\mathrm{net}}\,\Gamma_{k\ell,t},$$

where

$$\Pi_{ij,t}^{\mathrm{net}}\ \equiv\ R_{ij}[\pi(\mu_i^*) - \pi(\tilde{\mu}_{ij,t})].$$

Note that $\pi(\mu_i^*) - \pi(\tilde{\mu}_{ij,t}) > 0$ (per-period profit gain at frictionless elasticity, first order in $\omega_{ij,t}$ since μ_i^* maximizes $\pi(\mu)$). At $\omega_{ij,t} = 0$: $\tilde{\mu}_{ij,t} = \mu_i^*$, so $\Pi_{ij,t}^{\mathrm{net}} = 0$. For $\omega_{ij,t} > 0$, using $\omega_{ij,t} = O(\sigma_z^2/\sigma_{\zeta,j}^2)$:

$$\Pi_{ij,t}^{\mathrm{net}}\ =\ c_\Pi\,R_{ij}\,\sigma_z^2\ +\ O(\sigma_z^4), \tag{35}$$

where $c_\Pi \equiv \varepsilon_i|W'_{k\ell,t}(\mu_i^*)|/[\Gamma_{k\ell,t}(\varepsilon_i-1)^2\sigma_{\zeta,j}^2] > 0$. Substituting and dividing by $\Gamma_{k\ell,t} > 0$, the IC becomes

$$\psi\ \le\ c_\Pi\,R_{ij}\,\sigma_{z,j}^2\,\Gamma_{k\ell,t}\ +\ O(\sigma_z^4).$$

Since $R_{ij} = 1/(1-\beta\rho_{ij})$ is increasing in ρ_{ij}, the IC binds at a unique threshold $\bar{\rho}_{ij} \in (0,1)$ satisfying $\psi = c_\Pi\,\bar{R}_{ij}\,\sigma_{z,j}^2\,\Gamma_{k\ell,t}$, where $\bar{R}_{ij} \equiv 1/(1-\beta\bar{\rho}_{ij})$. To leading order in $\sigma_{z,j}^2$, inverting gives:

$$\bar{\rho}_{ij}\ \approx\ 1 - \frac{1}{\beta} + \frac{\psi}{\beta\,c_\Pi\,\sigma_{z,j}^2\,\Gamma_{k\ell,t}}. \tag{36}$$

Step 3: Comparative statics. From (36):

(a) *Decreasing in $\sigma_{z,j}^2$.* The threshold $\bar{\rho}_{ij}$ is decreasing in $\sigma_{z,j}^2$ because higher aggregate volatility raises ω_{ij} and hence the markup distortion $\tilde{\mu}_{ij} - \mu_i^*$, enlarging the gains from disclosure.

(b) *Increasing in $\sigma_{\zeta,j}^2/\sigma_{z,j}^2$.* Higher $\sigma_{\zeta,j}^2$ reduces c_Π through the denominator, raising $\bar{\rho}_{ij}$. Intuitively, noisier bilateral prices reduce the Kalman weight ω_{ij}, shrinking the markup distortion and the gains from disclosure.

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

The discount factor β enters R_{ij} multiplicatively: higher β lowers $\bar{\rho}_{ij}$ uniformly, 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 (14) 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, $P_{j,t}$ 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

$$x^{\text{int}}(p) \;=\; \frac{a_{ij}\, S_{i,t}}{p}, \tag{37}$$

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 $P_{j,t}$. Switching entails a match-specific cost $\tau\, P_{j,t}$, 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 $P_{j,t}$ directly. As established in the main text, $P_{j,t} = \bar{P}_j\, e^{z_{j,t}}$, so uncertainty about $P_{j,t}$ is equivalent to uncertainty about the common component $z_{j,t}$. The buyer evaluates the outside option at the conditional expectation $\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$.

The buyer sources from supplier ℓ if

$$p \;\leq\; \tau\, \mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]\,. \tag{38}$$

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

$$\Pr(\text{source from } \ell) = 1 - G\!\left(\frac{p}{\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]}\right). \tag{39}$$

D.3 Aggregate bilateral demand

Combining the intensive margin (37) with the extensive margin (39), the bilateral demand at price p is

$$x(p) = \frac{a_{ij}\, S_{i,t}}{p}\left[1 - G\!\left(\frac{p}{\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]}\right)\right]. \tag{40}$$

D.4 CES specification

The bilateral demand (40) has constant price elasticity ε_i if and only if the survival function $1 - G$ is a power law:

$$1 - G(\tau) = \tau^{-(\varepsilon_i - 1)}, \qquad \tau \geq 1,\ \varepsilon_i > 1. \tag{41}$$

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 (41), equation (40) reduces to the CES form.

Derivation. Let $r \equiv p/\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$. Substituting (41) into (40):

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

□

This yields the CES demand

$$x_{kj,t}(p) \;=\; \left(\frac{p}{\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]}\right)^{-\varepsilon_i} A_{kj,t}, \qquad A_{kj,t} \;\equiv\; \frac{a_{ij}\, S_{i,t}}{\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]}, \tag{42}$$

which is equation (14). 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 (41); the intensive margin contributes one unit from (37).

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

On *disclosure links*, the buyer observes the supplier's belief $\hat{z}_{j,t}$ directly, so $\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$ 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 $z_{j,t}$, 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 (20) follows.

E 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 $\mathcal{R}_t$.

E.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 firm's own perfectly observed $\xi_{k,t}$ and is held fixed thereafter. The signal extraction problem below therefore conditions on $\hat{\zeta}_{ik,t}$ and concerns only the common component $z_{i,t}$.

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

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

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

(c) for each non-disclosure link $(j,i) \notin \mathcal{R}_t$, the price residual $\hat{z}_{j,t} + \hat{\zeta}_{j\ell_j(k),t}$ with Kalman weight $\omega_{ij,t}$ from equation (17).

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

$$\hat{z}_{ik,t} = \kappa_i \, s_{z,k,i,t} \; + \sum_{j:\,(j,i)\in\mathcal{R}_t} \psi_{ij,t} \, \hat{z}_{j,t} \; + \sum_{j:\,(j,i)\notin\mathcal{R}_t} \beta^p_{ij,t} \left(\hat{z}_{j,t} + \hat{\zeta}_{j\ell_j(k),t} \right), \tag{43}$$

where the weights $(\kappa_i, \psi_{ij,t}, \beta^p_{ij,t})$ are the partial regression coefficients from the joint regression of $z_{i,t}$ 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 $z_{j,t}$ is highly correlated with $z_{j',t}$ 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 (17), which governs the markup distortion.[43] On disclosure links, sufficiency of the disclosed belief implies $\beta^p_{ij,t} = 0$: the buyer ignores the price residual for that supplier.

E.2 Aggregation to industry-level beliefs

Integrating (43) over $k \in [0,1]$ and applying the law of large numbers (the idiosyncratic noise $\hat{\zeta}_{j\ell_j(k),t}$ and $\nu_{k,t}$ are independent across the different partners that different firms k face), the industry-average belief satisfies equation (44) in the main text:

$$\bar{\hat{z}}_{i,t} = \kappa_i \, z_{i,t} + \sum_{j:\,(j,i)\in\mathcal{R}_t} \psi_{ij,t} \, \bar{\hat{z}}_{j,t} + \sum_{j:\,(j,i)\notin\mathcal{R}_t} \beta^p_{ij,t} \, z_{j,t} + \eta_{i,t}, \tag{44}$$

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

E.3 Matrix characterization

Stacking across n industries, define:

- Ψ^R_t: the $n \times n$ matrix with (i,j)-entry $\psi_{ij,t} \, \mathbf{1}[(j,i) \in \mathcal{R}_t]$ (disclosed-belief weights, upstream);

[43] 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. The belief weight $\beta^p_{ij,t}$ reflects the price signal's contribution to the full posterior, after controlling for all other signals.

- $\kappa \equiv \text{diag}(\kappa_1, \ldots, \kappa_n)$ (private-signal weights);
- B_t^{NR}: the $n \times n$ matrix with (i,j)-entry $\beta_{ij,t}^{p}\, \phi_j\, \mathbf{1}[(j,i) \notin \mathcal{R}_t]$ (price-channel contribution to exogenous loading on $z_{j,t}$).

Define $B_t^{\text{exog}} \equiv \kappa + B_t^{NR}$ (the exogenous block: weights on fundamental $\mathbf{z}_t$) and $B_t^{\text{endo}} \equiv \Psi_t^R$ (the endogenous block: weights on other industries' beliefs). The system (45) is

$$\bar{\hat{\mathbf{z}}}_t = B_t^{\text{endo}}\, \bar{\hat{\mathbf{z}}}_t + B_t^{\text{exog}}\, \mathbf{z}_t + \boldsymbol{\eta}_t. \tag{45}$$

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 (46):

$$\bar{\hat{\mathbf{z}}}_t = [I - B_t^{\text{endo}}]^{-1}\, B_t^{\text{exog}}\, \mathbf{z}_t + [I - B_t^{\text{endo}}]^{-1} \boldsymbol{\eta}_t. \tag{46}$$

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 $\mathbf{z}_t$ into beliefs $\bar{\hat{\mathbf{z}}}_t$; its diagonal entry $M_{ii,t}$ is the accuracy of industry i's average belief about its own common component.

E.4 Disclosure set and fixed point

The disclosure set $\mathcal{R}_t$ is defined as the largest set of links satisfying the establishment condition (24) simultaneously:

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

where $R_{ij} \equiv 1/(1 - \beta\rho_{ij})$ capitalizes per-period flows over the expected match life, ψ is the uniform setup cost (see Section 5.3), and $\Delta W_{ji,t}$ is the cooperative surplus (22). The cooperative surplus depends on the bilateral markup weight $\omega_{ij,t}$ from (17), which in turn depends on the disclosure configuration $\mathcal{R}_t$. The equilibrium is therefore a fixed point.

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 $\mathcal{R}_t^{(s+1)} = \{(j,i) : f \le R_{ij}\, \Delta W_{ji,t}^{(s)}\}$.

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.

F 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:

$$\mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\log C_t - \frac{L_t^{1+\varphi}}{1+\varphi} \right], \tag{47}$$

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:

$$C_t = \prod_{i=1}^{n} C_{i,t}^{\theta_i}, \qquad \sum_{i=1}^{n} \theta_i = 1, \tag{48}$$

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.[44]

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

$$L_t^{\varphi} = \frac{w_t}{P_t}. \tag{49}$$

Monetary policy

Monetary policy sets the nominal wage w_t directly. This is a standard modeling device in the dispersed-information literature.[45] The policy rule is

$$\log w_t = \bar{w} + \varsigma_t, \tag{50}$$

[44] 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.

[45] 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.

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

Under (50), the common component $z_{i,t}$ defined in (10) becomes

$$z_{i,t} = \alpha_i \, \varsigma_t + \epsilon_{i,t} + \sum_{j=1}^{n} a_{ij} \, z_{j,t}, \tag{51}$$

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

Production technology

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

$$y_{ik,t} = e^{\epsilon_{i,t} + \xi_{k,t}} \, l_{ik,t}^{\alpha_i} \prod_{j=1}^{n} m_{ijk,t}^{a_{ij}}, \tag{52}$$

where $l_{ik,t}$ is labor, $m_{ijk,t}$ is the intermediate input from industry j, $\epsilon_{i,t}$ is an industry-common productivity shock, and $\xi_{k,t}$ is a firm-specific productivity shock.

Market clearing

Labor market. The labor market clears at w_t:

$$L_t = \sum_{i=1}^{n} \int_0^1 l_{ik,t} \, dk. \tag{53}$$

From the Cobb–Douglas technology (52), 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.

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

$$Y_{i,t} = C_{i,t} + \sum_{j=1}^{n} \int_0^1 x_{ij,k,t} \, dk. \tag{54}$$

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 (13), the industry-level log price is $\log P_{i,t} = \bar{P}_i + \bar{\hat{z}}_{i,t}$ (as derived in (57)). The aggregate price index is

$$\log P_t = \sum_{i=1}^{n} \theta_i \log P_{i,t} = \bar{P} + \sum_{i=1}^{n} \theta_i \, \bar{\hat{z}}_{i,t}. \tag{55}$$

Combining with the labor supply condition (49):

$$\varphi \, \log L_t = \log w_t - \log P_t = \varsigma_t + \bar{w} - \bar{P} - \sum_{i=1}^{n} \theta_i \, \bar{\hat{z}}_{i,t}, \tag{56}$$

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 $\bar{\hat{z}}_{i,t}$ are determined by (46), 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 $\mathbf{z}_t$ and the signal structure, but not on quantities or the aggregate price level.[46]

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.

Efficient benchmark. The efficient benchmark is the *constrained efficient allocation*: all firms observe $z_{i,t}$ perfectly ($\bar{\hat{z}}_{i,t} = z_{i,t}$ 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.

[46] 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.

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

$$\log P_{i,t} = \overline{\log p_{ik,t}} - \frac{\varepsilon_i - 1}{2} \operatorname{Var}_k(\log p_{ik,t}) \,, \tag{57}$$

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

$$m_i \equiv \sum_{j=1}^{n} a_{ij}(1 - r_t^{ji})\, \delta_{ij,t}, \qquad \delta_{ij,t} \equiv \log \tilde{\mu}_{ij,t} - \log \mu_i^* = \log \frac{\varepsilon_i(1-\omega_{ij,t})}{\varepsilon_i(1-\omega_{ij,t}) - 1} - \log \frac{\varepsilon_i}{\varepsilon_i - 1}, \tag{58}$$

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 (20) 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 $\bar{\hat{z}}_{i,t} = z_{i,t}$, so $\log P_{i,t}^{**} = \bar{P}_i + z_{i,t} - \frac{\varepsilon_i - 1}{2} V_i^*$, 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

$$\log P_t - \log P_t^{**} = \underbrace{\bar{m}}_{\text{markup level}} - \underbrace{\tilde{e}_t}_{\text{belief error}} - \underbrace{\frac{1}{2} \sum_{i=1}^{n} \theta_i (\varepsilon_i - 1) \Delta V_{i,t}}_{\text{dispersion change}}, \tag{59}$$

where $\bar{m} \equiv \sum_i \theta_i m_i$, $\tilde{e}_t \equiv \sum_i \theta_i (z_{i,t} - \bar{\hat{z}}_{i,t})$, 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 $s_{ik,t} = z_{i,t} + \nu_{ik,t}$ only. The Kalman gain $\phi_i = \sigma_{z,i}^2/(\sigma_{z,i}^2 + \sigma_{\nu,i}^2)$ gives the industry-average log price as $\overline{\log p_{ik,t}} = \bar{P}_i + m_i + \phi_i z_{i,t}$, and each firm's deviation from the industry average is:

$$\log p_{ik,t} - \overline{\log p_{ik,t}} = \phi_i \, \nu_{ik,t} + (\hat{\zeta}_{ik,t} - \bar{\hat{\zeta}}_{i,t}), \tag{60}$$

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 z_i and present even under full information). The excess within-industry price dispersion due to imperfect information is therefore:

$$\Delta V_{i,t} = \phi_i^2 \, \sigma_{\nu,i}^2. \tag{61}$$

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 (61) 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 z_i that enters the aggregate price level in (59) is $\bar{\hat{z}}_{i,t} = \phi_i z_{i,t}$ (the own-signal posterior mean), so the aggregate belief error $\tilde{e}_t = \sum_i \theta_i (1 - \phi_i) z_{i,t}$ depends only on ϕ_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 (47) and (49). 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:

$$U'(x) = \frac{1+\varphi}{\varphi} - \frac{1}{\varphi} e^{x(1+\varphi)/\varphi}, \qquad U''(x) = -\frac{(1+\varphi)^2}{\varphi^2} e^{x(1+\varphi)/\varphi}. \tag{62}$$

To second order around x^{**}:

$$\mathcal{L} \equiv \mathbf{E}[U(x^{**}) - U(x_t)] \approx -U'(x^{**})\,\mathbf{E}[x_t - x^{**}] - \tfrac{1}{2}U''(x^{**})\,\mathbf{E}[(x_t - x^{**})^2]. \tag{63}$$

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 (59):

$$\mathbf{E}[x_t - x^{**}] = -\mathbf{E}[\log P_t - \log P^{**}] = -\bar{m} + O(\omega^2), \tag{64}$$

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 $-U''(x^{**}) = (1+\varphi)^2 e^{x^{**}(1+\varphi)/\varphi}/\varphi^2$:

$$\mathbf{E}[(x_t - x^{**})^2] = \mathrm{Var}(\tilde{e}_t) + O(\omega^3), \tag{65}$$

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.

$$\mathcal{L} = \underbrace{U'(x^{**})\,\bar{m}}_{\substack{\mathcal{W}_1\text{: markup cost of non-disclosure} \\ \text{(first order in } \omega\text{)}}} + \underbrace{\frac{(1+\varphi)^2}{2\varphi^2} e^{x^{**}(1+\varphi)/\varphi}\ \mathrm{Var}(\tilde{e}_t)}_{\substack{\mathcal{W}_2\text{: aggregate belief error} \\ \text{(La'O first term, second order)}}} + \underbrace{\frac{1}{2}\sum_{i=1}^{n} \theta_i(\varepsilon_i - 1)\,\mathbf{E}[\Delta V_{i,t}]}_{\substack{\mathcal{W}_3\text{: within-industry belief dispersion} \\ \text{(La'O second term, second order)}}}, \tag{66}$$

where $\tilde{e}_t = \sum_i \theta_i(z_{i,t} - \bar{\hat{z}}_{i,t})$, $\bar{m} = \sum_i \theta_i m_i$, $\Delta V_{i,t}$ is given by (61), 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.[47]

[47]The formula captures welfare costs arising from seller-side belief distortions—the markup wedge on non-disclosure links ($\mathcal{W}_1$) and the seller pricing errors induced by imperfect information ($\mathcal{W}_2$, $\mathcal{W}_3$)—but abstracts from a fourth channel: the welfare cost of buyer switching decision errors. Buyers decide whether to switch to the wholesale option based on $\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$ rather than the true $P_{j,t}$; when beliefs are incorrect, some buyers switch who should not (paying iceberg waste $(\tau - 1)P_{j,t}$ 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 $P_{j,t}$, 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 $\mathcal{W}_1$. 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 $\mathcal{W}_2$ and $\mathcal{W}_3$ are slightly understated; the overall welfare cost (66) is therefore a conservative lower bound on the full second-order cost.

Interpreting the three terms. $\mathcal{W}_1$ 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. $\mathcal{W}_1$ 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.

$\mathcal{W}_1$ 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 (22). 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 $\mathcal{W}_1$ uses consumption weights $\theta_i a_{ij}$ through $\bar{m}$. The ratio $\mathrm{T1}/\mathcal{L}_\mu = U'(x^{**}) \cdot \sum_i \theta_i m_i / \sum_{(i,j)} \lambda_i(1-\alpha_i)a_{ij}|W'_{ij}|\delta_{ij}$ is a model-specific constant determined by the steady-state network aggregation. We compute $\mathcal{W}_1$ directly from the GE formula (66) for Table 11, and report $\mathcal{L}_\mu$ as a supplementary quantity in Section 5.3.

$\mathcal{W}_2$ 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, $\tilde{e}_t = \sum_i \theta_i(1-\phi_i)z_{i,t}$ and $\mathrm{Var}(\tilde{e}_t)$ depends only on ϕ_i and Σ_z. $\mathcal{W}_2$ 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.

$\mathcal{W}_3$ 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 (61), $\Delta V_{i,t} = \phi_i^2 \sigma_{\nu,i}^2$ depends only on the seller's own-signal precision. Like $\mathcal{W}_2$, $\mathcal{W}_3$ is *constant across disclosure regimes*. It vanishes under full information (all firms observe z_i perfectly) but does not respond to changes in the bilateral learning structure.

Online Appendix

(not for publication)

G Surplus derivations and breadth–noise condition

Bilateral surplus derivations

Using the CES demand (14), the supplier's per-period profit on the link between buyer k and supplier ℓ is

$$\Pi^S_{k\ell,t}(\mu) = (p - mc_{j\ell,t})\, x_{kj,t} = (\mu - 1)\mu^{-\varepsilon}\Gamma_{k\ell,t} \equiv \pi(\mu)\,\Gamma_{k\ell,t}. \tag{67}$$

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 $\Gamma_{k\ell,t} \equiv mc^{1-\varepsilon}_{j\ell,t} A_{kj,t}$ gives

$$\Pi^B_{k\ell,t}(\mu) = \frac{1}{\varepsilon - 1}\,\mu^{1-\varepsilon}\,\Gamma_{k\ell,t}. \tag{68}$$

Total bilateral surplus is

$$W_{k\ell,t}(\mu) \equiv \Pi^S_{k\ell,t}(\mu) + \Pi^B_{k\ell,t}(\mu) = \left[(\mu - 1)\mu^{-\varepsilon} + \frac{\mu^{1-\varepsilon}}{\varepsilon - 1}\right]\Gamma_{k\ell,t}. \tag{69}$$

Differentiating:

$$W'_{k\ell,t}(\mu) = \varepsilon\,\mu^{-\varepsilon-1}(1 - \mu)\,\Gamma_{k\ell,t}. \tag{70}$$

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

$$\pi'(\mu^*_i) = 0, \qquad \pi''(\mu^*_i) < 0. \tag{71}$$

Breadth–noise scaling

Supplier breadth has *competing* effects on $\omega_{ij,t}$ (17). 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 $z_{j,t} = \sum_r h_{jr}\alpha_r(\log w_t - \bar{w}) + \sum_r h_{jr}e_{r,t}$ from (10), which loads on many diverse industry-specific shocks that the buyer cannot easily learn from its private signal $s_{z,k,i,t}$ or from the upstream attention signal. The net effect is that $\omega_{ij,t}$ increases with supplier breadth under the following condition.

Assumption 1 (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 $z_{j,t}$ through the Leontief inverse). The industry-specific shocks $\{e_{r,t}\}_{r=1}^n$ have uniformly bounded variances $\underline{\sigma}^2_e \leq \mathrm{Var}(e_{r,t}) \leq \bar{\sigma}^2_e$. Then for all j and all buyers k with supplier ℓ in industry j:*

$$\gamma_{ij,t} \equiv \mathrm{Var}(\hat{z}_{j\ell,t} \mid s_{z,k,i,t},\, m_{k\ell,t}) \;\geq\; c\, B_j\, \underline{\sigma}^2_e \qquad \textit{for some } c > 0.$$

Assumption 1 holds when the Leontief weights h_{jr} spread across many industries with independent shocks $e_{r,t}$: the common component $z_{j,t} = \sum_r h_{jr}e_{r,t} + \cdots$ becomes a weighted average of many independent terms, each of which is hard for the buyer to learn individually from its private signal about $z_{i,t}$ (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 1, the numerator of (17) 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.[48]

H 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 $z_{j,t}$ are the private signal $s_{z,k,i,t}$ and the bilateral price residual (15). 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.

H.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:

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

where $\lambda_j > 0$ measures the sensitivity of industry j's activity to beliefs and $u_{j\ell,t}$ is independent of $(z_{j,t}, \zeta_{j\ell,t})$. 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

$$s^{\text{up}}_{k\ell,t} \equiv \frac{a_{j\ell,t} + \nu^{\text{up}}_{k\ell,t}}{\lambda_j} = \hat{z}_{j\ell,t} + \underbrace{\frac{u_{j\ell,t} + \nu^{\text{up}}_{k\ell,t}}{\lambda_j}}_{\equiv\ \eta^{\text{up}}_{k\ell,t}}, \tag{72}$$

[48] When non-cooperative signals are introduced (Appendix E and Appendix H), the assumption acquires an additional term $-\iota^{-1}_{\text{up},k\ell}$ 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.

with noise variance

$$\sigma^2_{\text{up},k\ell} = \frac{\sigma^2_{u,j} + \sigma^2_{\nu,k\ell}}{\lambda_j^2}. \tag{73}$$

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 (15) are conditionally independent given $\hat{z}_{j\ell,t}$. 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 $\mathcal{I}^{\text{bi}}_{k\ell,t} = \{s_{z,k,i,t},\, s^{\text{up}}_{k\ell,t},\, \text{price residual}\}$, and the residual uncertainty $\gamma_{ij,t}$ in (17) 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 (13), each of buyer k's output prices has the form

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

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 $\hat{z}_{ik,t} + \hat{\zeta}_{ik,t}$.

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

$$s^{\text{dn}}_{\ell k,t} = \hat{z}_{ik,t} + \underbrace{\hat{\zeta}_{ik,t} + \nu^{\text{dn}}_{\ell k,t}}_{\equiv\ \eta^{\text{dn}}_{\ell k,t}}, \tag{74}$$

with noise variance

$$\sigma^2_{\text{dn},\ell k} = \sigma^2_{\zeta,i} + \sigma^2_{\nu,\ell k}, \tag{75}$$

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 $z_{j,t}$ because $z_{i,t}$ and $z_{j,t}$ 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 (15) given $\hat{z}_{ik,t}$ and $\hat{z}_{j\ell,t}$.

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 $z_{j,t}$), 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.

H.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:

$$\sigma^2_{\nu,k\ell} = \frac{\sigma^2_0}{\lambda^{\text{att}}_{k\ell}}, \tag{76}$$

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 $c'' > 0$.[49]

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 $\sigma^2_{z,j}$. Counterfactual exercises consider unforeseen (MIT) shocks to $\sigma^2_{z,j}$, after which firms re-optimize attention and disclosure given the new level.[50]

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

(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:

$$\left(\{\lambda^{\text{att},*}_{k\ell,t}\},\ \mathcal{R}_t\right) \quad \text{such that} \quad \begin{cases} \lambda^{\text{att},*}_{k\ell,t} \text{ is optimal given } \mathcal{R}_t \text{ and } \sigma^2_{z,j}, \\ \mathcal{R}_t \text{ satisfies (24) given } \sigma^2_{\nu,k\ell}(\lambda^{\text{att},*}_{k\ell,t}). \end{cases}$$

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

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

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

$$\max_{\lambda_{k\ell}^{\text{att}} \geq 0} \; a_{ij}\, \Gamma_{k\ell,t}\, g\Big(\omega_{ij,t}(\lambda_{k\ell}^{\text{att}}, \sigma_{z,j}^2)\Big) \;-\; c\Big(\lambda_{k\ell}^{\text{att}}\Big), \tag{77}$$

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:

$$a_{ij}\, \Gamma_{k\ell,t}\, g'(\omega)\, \frac{\partial \omega}{\partial \lambda_{k\ell}^{\text{att}}} \;=\; c'(\lambda_{k\ell,t}^{\text{att},*}). \tag{78}$$

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

$$\lambda_{k\ell,t}^{\text{att},*} \;=\; \frac{1}{\kappa_i^{\text{att}}}\, a_{ij}\, \Gamma_{k\ell,t}\, \mathcal{V}_{k\ell,t}(\sigma_{z,j}^2), \tag{79}$$

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

Properties of optimal attention. Four properties of (79) 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 $\sigma_{z,j}^2$: when aggregate uncertainty is high, the markup distortion is more costly and the marginal value of attention rises. Following an MIT shock that raises $\sigma_{z,j}^2$, *both* margins activate simultaneously: disclosure links expand (Proposition 3) *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 $\hat{z}_{j\ell,t}$ 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 2(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 3. With non-cooperative signals, the cooperative surplus coefficient c_W in the proof of Proposition 3 depends on $\sigma_{\text{up},k\ell}^2$: 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 $\ell_{k\ell}^{(k)}$ (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.

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.[51] 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.

[51]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 Horizon-specific learning estimates

This appendix reports the joint-direction learning regression at the two individual forecast horizons, $h = 6$ and $h = 12$ quarters ahead. The specification mirrors the main table (Table 3) exactly, except that the horizon is fixed rather than pooled. The column structure, weighting schemes, fixed effects, and inference procedures are identical to the main table.

Main patterns. The qualitative structure of the main results is stable across horizons. The downstream belief channel ($T_j \times$ expect_down) is positive and significant at both horizons, and is substantially stronger at $h = 12$ than at $h = 6$—consistent with longer-horizon expectations being more informative about underlying cost conditions. The upstream belief channel is weaker at $h = 6$ but becomes more precisely estimated at $h = 12$, particularly under transaction-value and breadth weights. The price channel upstream shows stronger evidence at $h = 12$, with significant coefficients under breadth and transaction-value weights (0.017** and 0.019**).

Tables L.3 and L.4 report the same horizon $h = 12$ analysis estimated separately by direction (suppliers only; customers only), as a robustness check that the joint specification does not drive the results. The qualitative patterns are unchanged: the upstream belief channel is significant under imp and impJ weights ($\delta_1 = 0.080^{*}$ and 0.053^{**}); the downstream belief channel is uniformly large and highly significant across all weighting schemes ($\delta_1 = 0.088^{***}$ to 0.157^{***}).

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

Table L.1: Learning from supply-chain partners ($h = 6$)

	Signal only	Horse race (UNT only)				Triple interactions (ALL)			
	(1) unif	(2) unif	(3) supA	(4) imp	(5) impJ	(6) unif	(7) supA	(8) imp	(9) impJ
price_up	0.002	0.001	-0.000	0.001	0.010	-0.001	-0.003	-0.001	0.011
	(0.003)	(0.004)	(0.003)	(0.004)	(0.008)	(0.004)	(0.004)	(0.005)	(0.008)
1.j_treat#c.price_up	0.000	0.008	0.007	-0.002	-0.003	0.009	0.010*	0.000	-0.004
	(0.004)	(0.007)	(0.005)	(0.004)	(0.009)	(0.007)	(0.006)	(0.005)	(0.008)
price_down	-0.001	-0.002	-0.001	-0.000	-0.001	-0.001	-0.001	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
1.j_treat#c.price_down	0.002	0.002	0.001	0.000	0.002	0.002	0.001	-0.000	0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)
y_lag	0.258***	0.247***	0.258***	0.257***	0.257***	0.252***	0.259***	0.259***	0.259***
	(0.054)	(0.062)	(0.054)	(0.054)	(0.054)	(0.048)	(0.047)	(0.047)	(0.047)
expect_up		-0.004	-0.002	-0.012	0.012	-0.007	0.001	-0.018	0.028
		(0.015)	(0.015)	(0.020)	(0.024)	(0.022)	(0.021)	(0.023)	(0.024)
1.j_treat#c.expect_up		0.010	0.006	0.032	0.012	0.019	0.005	0.041	0.001
		(0.013)	(0.021)	(0.027)	(0.029)	(0.018)	(0.018)	(0.029)	(0.026)
expect_down		-0.023	-0.024	-0.034	0.001	-0.010	-0.017	-0.015	-0.010
		(0.019)	(0.019)	(0.022)	(0.023)	(0.018)	(0.023)	(0.018)	(0.018)
1.j_treat#c.expect_down		0.054***	0.054**	0.073**	0.035	0.031**	0.057*	0.051*	0.035
		(0.019)	(0.023)	(0.026)	(0.024)	(0.013)	(0.027)	(0.026)	(0.028)
1.firm_treat#c.price_up						0.002	0.003	0.003	-0.008
						(0.005)	(0.004)	(0.005)	(0.009)
1.j_treat#1.firm_treat#c.price_up						-0.008	-0.010*	0.007	0.007
						(0.008)	(0.005)	(0.005)	(0.008)
1.firm_treat#c.price_down						0.000	0.000	-0.001	-0.001
						(0.001)	(0.002)	(0.002)	(0.001)
1.j_treat#1.firm_treat#c.price_down						0.001	0.002	0.003	0.001
						(0.001)	(0.003)	(0.003)	(0.002)
1.firm_treat#c.expect_up						0.013	0.018	0.021	-0.029
						(0.023)	(0.021)	(0.021)	(0.024)
1.j_treat#1.firm_treat#c.expect_up						-0.021	-0.036	-0.032	0.014
						(0.015)	(0.021)	(0.024)	(0.023)
1.firm_treat#c.expect_down						0.026	0.048*	0.040*	0.032
						(0.016)	(0.025)	(0.023)	(0.026)
1.j_treat#1.firm_treat#c.expect_down						-0.027	-0.057	-0.028	-0.027
						(0.018)	(0.040)	(0.041)	(0.038)
Observations	83169	83169	83169	83169	83169	248649	248649	248649	248649
R-squared	0.065	0.060	0.066	0.065	0.066	0.062	0.065	0.065	0.065

Notes: The table reports the joint-direction pooled-horizon regression in equation (6), estimated on the h-quarter-ahead horizon indicated in the caption. Both upstream (supplier) and downstream (customer) signals appear jointly. Column (1) reports the transaction-only specification. Columns (2)–(5) add the partner belief signal for untreated firms only; columns (6)–(9) add own-treatment interactions. Weighting schemes: *unif* (uniform), *supA* (breadth-weighted), *imp* (transaction-value weighted), *impJ* (relationship-intensity weighted). All weighting scores are computed over the full sample and are time-invariant. All specifications include firm, sector×quarter, size×quarter, area×quarter fixed effects, and direction×treatment-pair×quarter fixed effects (columns 6–9 only). 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table L.2: Learning from supply-chain partners ($h = 12$)

	Signal only	Horse race (UNT only)				Triple interactions (ALL)			
	(1) unif	(2) unif	(3) supA	(4) imp	(5) impJ	(6) unif	(7) supA	(8) imp	(9) impJ
price_up	-0.004	-0.006	-0.009	-0.009	0.003	-0.008	-0.010*	-0.011	0.004
	(0.004)	(0.005)	(0.006)	(0.006)	(0.008)	(0.006)	(0.005)	(0.007)	(0.008)
1.j_treat#c.price_up	0.008	0.014	0.017**	0.007	0.002	0.017*	0.019**	0.010	0.002
	(0.005)	(0.008)	(0.007)	(0.006)	(0.008)	(0.009)	(0.007)	(0.007)	(0.008)
price_down	-0.001	-0.001	-0.000	0.001	0.002	0.001	0.001	0.002	0.004*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
1.j_treat#c.price_down	0.001	0.001	0.001	-0.001	-0.001	-0.000	-0.001	-0.003	-0.003
	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
y_lag	0.213***	0.204***	0.213***	0.211***	0.211***	0.252***	0.260***	0.259***	0.259***
	(0.063)	(0.070)	(0.062)	(0.063)	(0.063)	(0.052)	(0.052)	(0.052)	(0.052)
expect_up		-0.029	-0.031	-0.043	-0.035	-0.034	-0.016	-0.044	-0.020
		(0.021)	(0.021)	(0.034)	(0.037)	(0.030)	(0.021)	(0.036)	(0.039)
1.j_treat#c.expect_up		0.026	0.036	0.073*	0.059	0.032*	0.006	0.072*	0.048
		(0.015)	(0.029)	(0.037)	(0.035)	(0.019)	(0.020)	(0.035)	(0.032)
expect_down		-0.070**	-0.077**	-0.098**	-0.062	-0.058***	-0.094***	-0.089***	-0.066*
		(0.028)	(0.029)	(0.041)	(0.039)	(0.018)	(0.029)	(0.031)	(0.034)
1.j_treat#c.expect_down		0.096***	0.107***	0.148***	0.095**	0.078***	0.143***	0.130***	0.087*
		(0.025)	(0.031)	(0.042)	(0.039)	(0.024)	(0.040)	(0.042)	(0.045)
1.firm_treat#c.price_up						0.008	0.010	0.013*	-0.003
						(0.007)	(0.006)	(0.007)	(0.008)
1.j_treat#1.firm_treat#c.price_up						-0.014	-0.019**	-0.005	0.002
						(0.011)	(0.007)	(0.007)	(0.008)
1.firm_treat#c.price_down						-0.001	-0.001	-0.002*	-0.004**
						(0.001)	(0.001)	(0.001)	(0.002)
1.j_treat#1.firm_treat#c.price_down						0.002	0.003	0.005*	0.005*
						(0.001)	(0.003)	(0.003)	(0.003)
1.firm_treat#c.expect_up						0.041	0.025	0.049	0.022
						(0.037)	(0.027)	(0.038)	(0.046)
1.j_treat#1.firm_treat#c.expect_up						-0.045**	-0.024	-0.060**	-0.041*
						(0.021)	(0.023)	(0.024)	(0.023)
1.firm_treat#c.expect_down						0.089***	0.122***	0.110**	0.104**
						(0.027)	(0.042)	(0.040)	(0.045)
1.j_treat#1.firm_treat#c.expect_down						-0.106***	-0.151***	-0.108**	-0.112**
						(0.032)	(0.048)	(0.044)	(0.049)
Observations	83169	83169	83169	83169	83169	248649	248649	248649	248649
R-squared	0.044	0.042	0.046	0.047	0.045	0.061	0.064	0.065	0.064

Notes: The table reports the joint-direction pooled-horizon regression in equation (6), estimated on the h-quarter-ahead horizon indicated in the caption. Both upstream (supplier) and downstream (customer) signals appear jointly. Column (1) reports the transaction-only specification. Columns (2)–(5) add the partner belief signal for untreated firms only; columns (6)–(9) add own-treatment interactions. Weighting schemes: *unif* (uniform), *supA* (breadth-weighted), *imp* (transaction-value weighted), *impJ* (relationship-intensity weighted). All weighting scores are computed over the full sample and are time-invariant. All specifications include firm, sector×quarter, size×quarter, area×quarter fixed effects, and direction×treatment-pair×quarter fixed effects (columns 6–9 only). 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table L.3: Learning from suppliers, separate-direction ($h = 12$)

	Trans. only	Horse race (UNT only)				Triple interactions (ALL)			
	(1) supA	(2) supA	(3) unif	(4) imp	(5) impJ	(6) supA	(7) unif	(8) imp	(9) impJ
price_j	-0.010**	-0.012**	-0.007	-0.012	0.003	-0.012**	-0.008	-0.013*	0.003
	(0.004)	(0.005)	(0.006)	(0.008)	(0.010)	(0.005)	(0.005)	(0.007)	(0.010)
1.j_treat#c.price_j	0.015**	0.019**	0.015*	0.010	0.004	0.019**	0.016*	0.012	0.003
	(0.006)	(0.007)	(0.009)	(0.007)	(0.007)	(0.007)	(0.009)	(0.008)	(0.008)
y_lag	0.201**	0.249***	0.188**	0.198**	0.199**	0.249***	0.245***	0.248***	0.247***
	(0.074)	(0.055)	(0.078)	(0.074)	(0.075)	(0.055)	(0.056)	(0.055)	(0.055)
1.firm_treat#c.price_j		0.010*				0.010*	0.008	0.014*	-0.006
		(0.006)				(0.006)	(0.007)	(0.007)	(0.009)
1.j_treat#1.firm_treat#c.price_j		-0.019**				-0.019**	-0.014	-0.005	0.005
		(0.007)				(0.007)	(0.011)	(0.009)	(0.007)
expect		-0.015	-0.030	-0.049	-0.031	-0.015	-0.034	-0.046	-0.025
		(0.020)	(0.025)	(0.042)	(0.034)	(0.020)	(0.028)	(0.037)	(0.034)
1.j_treat#c.expect		0.008	0.025*	0.080*	0.053**	0.008	0.033*	0.075*	0.050**
		(0.022)	(0.014)	(0.042)	(0.020)	(0.022)	(0.018)	(0.037)	(0.022)
1.firm_treat#c.expect		0.027				0.027	0.041	0.054	0.034
		(0.028)				(0.028)	(0.036)	(0.040)	(0.042)
1.j_treat#1.firm_treat#c.expect		-0.031				-0.031	-0.048**	-0.065**	-0.038*
		(0.026)				(0.026)	(0.020)	(0.027)	(0.020)
N	44817	140301	44855	44855	44855	140301	140451	140451	140451
Observations	0.041	0.060	0.035	0.041	0.040	0.060	0.058	0.061	0.060

Notes: Separate-direction regression estimated on supplier links only (supplier_j = 1), at horizon $h = 12$. The specification follows equation (6) restricted to upstream signals: $g_{j,t}$ is supplier j's year-on-year own-price growth; $\pi^{e,h}_{j,t}$ is supplier j's 12-quarter-ahead inflation expectation. Column (1) reports the transaction-only specification (supA weights, untreated only). Columns (2)–(5) add the partner belief signal under four weighting schemes (untreated firms only). Columns (6)–(9) add own-treatment interactions (full sample). Weighting schemes: *supA* (breadth-weighted), *unif* (uniform), *imp* (transaction-value), *impJ* (relationship-intensity). All weighting scores are time-invariant global averages. Standard errors use the Driscoll–Kraay correction (bandwidth 2); firm-level clustering is used as a fallback. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table L.4: Learning from customers, separate-direction ($h = 12$)

	Trans. only	Horse race (UNT only)				Triple interactions (ALL)			
	(1) supA	(2) supA	(3) unif	(4) imp	(5) impJ	(6) supA	(7) unif	(8) imp	(9) impJ
gr_j_imponibile_100	0.002	0.002	0.000	0.002	0.002	0.002	0.001	0.003*	0.003**
	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
1.j_treat#c.gr_j_imponibile_100	-0.001	-0.002	0.000	-0.003	-0.000	-0.002	0.000	-0.003*	-0.001
	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)	(0.002)	(0.001)
y_lag	0.245***	0.292***	0.271***	0.243***	0.242***	0.292***	0.289***	0.291***	0.292***
	(0.048)	(0.053)	(0.044)	(0.048)	(0.047)	(0.053)	(0.049)	(0.053)	(0.052)
1.firm_treat#c.gr_j_imponibile_100		-0.001				-0.001	0.000	-0.002	-0.003*
		(0.001)				(0.001)	(0.001)	(0.001)	(0.001)
1.j_treat#1.firm_treat#c.gr_j_imponibile_100		0.004*				0.004*	0.001	0.005**	0.003
		(0.002)				(0.002)	(0.001)	(0.002)	(0.002)
expect		-0.082**	-0.067***	-0.110***	-0.063*	-0.082**	-0.066***	-0.109***	-0.062*
		(0.031)	(0.016)	(0.031)	(0.031)	(0.031)	(0.017)	(0.032)	(0.034)
1.j_treat#c.expect		0.115***	0.088***	0.157***	0.097**	0.115***	0.084***	0.152***	0.089**
		(0.035)	(0.020)	(0.039)	(0.038)	(0.035)	(0.023)	(0.041)	(0.040)
1.firm_treat#c.expect		0.106***				0.106***	0.087***	0.121***	0.082*
		(0.035)				(0.035)	(0.021)	(0.039)	(0.041)
1.j_treat#1.firm_treat#c.expect		-0.121***				-0.121***	-0.105***	-0.118**	-0.093*
		(0.040)				(0.040)	(0.027)	(0.044)	(0.038)
N	38271	108031	38271	38271	38271	108031	108031	108031	108031
Observations	0.056	0.076	0.072	0.062	0.058	0.076	0.074	0.077	0.076

Notes: Separate-direction regression estimated on customer links only (supplier_j = 0), at horizon $h = 12$. The specification follows equation (restricted to downstream signals: $g_{j,t}$ is customer j's year-on-year purchase growth (winsorized); $\pi^{e,h}_{j,t}$ is customer j's 12-quarter-ahead inflati expectation. Column structure and weighting schemes as in Table L.3. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

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.5 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.

Table L.5: 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 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 by incentives to share information, whereas price-based learning operates more uniformly across relationships.

Implied marginal effects. Appendix Table L.6 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.6: 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.7 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.8, L.9, and L.10 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.7: 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, averaged over the full sample. *Partner breadth* is defined as the time-averaged number of upstream suppliers of firm j, 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.8: 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.9: 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.10: 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.11 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.11: 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*), averaged over the full sample and demeaned within the estimation sample. "Partner breadth" uses the time-averaged upstream degree of partner j; "own breadth" maps the same object to firm i. 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.12 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.12: 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.13 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.14, L.15, and L.16 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.17 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.13: 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.14: 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*), averaged over the full sample and demeaned within the estimation sample. "Partner breadth" uses the time-averaged upstream degree of partner j; "own breadth" maps the same object to firm i. 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.15: 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.16: 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.17: Learning from supply-chain partners: symmetric price signal

	Untreated only ($T_{i,t}=0$)				Full sample			
	(1) unif	(2) supA	(3) imp	(4) impJ	(5) unif	(6) supA	(7) imp	(8) impJ
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 (price 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 price 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.18 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.18: Robustness: partner-level controls

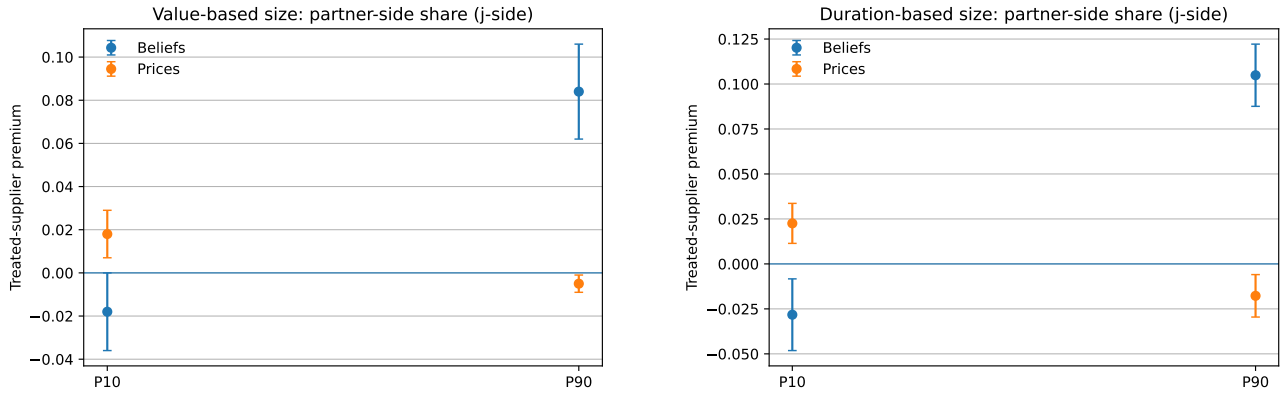
	Untreated only ($T_{i,t} = 0$)				Full sample			
	(1) brd	(2) unif	(3) val	(4) shr	(5) brd	(6) unif	(7) val	(8) shr
Panel A: Upstream signals ($g_{j,t}$ = supplier price growth $g^p_{j,t}$)								
β_0: $g^p_{j,t}$	−0.010*	−0.005	−0.005	0.003	−0.013*	−0.009	−0.009	0.004
	(0.005)	(0.005)	(0.005)	(0.009)	(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 (price 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 price 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 and marginally significant upstream belief interaction (0.033, s.e. 0.018), consistent with buyers exposed to more complex input environments extracting more from each supplier's belief signal. 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)	0.019	0.172**	0.153*
	(0.016)	(0.076)	(0.084)
$T_{j,t} \times g_{j,t}^{p}$	0.011	−0.041	−0.052
	(0.009)	(0.038)	(0.046)
$T_{j,t} \times \pi_{j,t}^{e,h}$ (downstream)	0.085***	0.016	−0.069*
	(0.024)	(0.023)	(0.039)
$T_{j,t} \times g_{j,t}^{r}$	0.001	0.005	0.003
	(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	−0.010	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: $T_{j,t}\, g_{j,t} \times \widetilde{b}_{i,t}^{(i)}$			−0.011	
			(0.010)	
δ_1: $T_{j,t}\, \pi^{e,h}_{j,t}$	0.028**	0.023	0.082**	0.035**
	(0.013)	(0.014)	(0.029)	(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: $T_{j,t}\, \pi^{e,h}_{j,t} \times \widetilde{b}_{i,t}^{(i)}$			0.033*	
			(0.018)	
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***	0.057***	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)	(0.060)	(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 time-averaged 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
$\pi_{j,t}^{e,h}$ (upstream)	0.583	0.752	0.489	0.613
$\pi_{j,t}^{e,h}$ (downstream)	0.219	0.410	0.091	0.129
$g_{j,t}^{p}$	0.583	0.307	0.715	0.898
$g_{j,t}^{r}$	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 (0.098) and the medium-breadth bin (0.057), but flips sign in the broad-breadth bin (−0.294), though the latter is estimated with substantial noise (s.e. 0.058). The downstream belief premium shows the opposite pattern: it is insignificant in the narrow bin (−0.024) but large and significant in the broad bin (0.179, s.e. 0.094).

Under own breadth (Panel B), the upstream belief premium is significant in the narrow bin (0.108, s.e. 0.053) and the medium bin (0.035**), but small and insignificant in the broad bin. The downstream belief premium is significant in the broad bin (0.014**).

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)
$T_{j,t} \times g^{p}_{j,t}$	0.015	−0.000	−0.033	−0.023	−0.010***	−0.020
	(0.044)	(0.008)	(0.036)	(0.023)	(0.003)	(0.021)
$\pi^{e,h}_{j,t}$ (upstream)	−0.047	−0.016	0.185**	0.067	−0.010	0.204*
	(0.061)	(0.017)	(0.081)	(0.046)	(0.020)	(0.113)
$T_{j,t} \times \pi^{e,h}_{j,t}$ (upstream)	0.098*	0.057***	−0.294***	0.108*	0.035**	−0.020
	(0.055)	(0.017)	(0.058)	(0.053)	(0.013)	(0.036)
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)
$T_{j,t} \times g^{r}_{j,t}$	−0.000	−0.002	0.011	0.007	0.004*	0.001
	(0.003)	(0.002)	(0.007)	(0.007)	(0.002)	(0.001)
$\pi^{e,h}_{j,t}$ (downstream)	0.027**	−0.004	−0.151	−0.015	0.034	0.024***
	(0.012)	(0.020)	(0.106)	(0.031)	(0.022)	(0.007)
$T_{j,t} \times \pi^{e,h}_{j,t}$ (downstream)	−0.024	0.026	0.179*	0.030	0.005	0.014**
	(0.023)	(0.021)	(0.094)	(0.020)	(0.018)	(0.006)
Lagged expectation $\pi^{e,h}_{i,t-1}$	0.293***	0.306***	0.315***	0.226***	0.349***	0.527***
	(0.026)	(0.040)	(0.061)	(0.066)	(0.042)	(0.071)
Observations	53,264	57,940	55,126	54,982	58,038	53,350
R^2	0.099	0.099	0.078	0.057	0.157	0.322

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 (0.016), with a marginally significant negative gradient ($\Delta = -0.069$, $p = 0.090$; 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:

$$
\begin{aligned}
\pi^e_{i,t} = \alpha_i + \alpha_t + \rho\,\pi^e_{i,t-1} + X_{ij,t}\varpi \\
&+ \beta_0 g_{j,t} + \beta_1(T_j \times g_{j,t}) + \beta_2(T_i \times g_{j,t}) + \beta_3(T_i \times T_j \times g_{j,t}) \\
&+ \delta_0 \pi^e_{j,t} + \delta_1(T_j \times \pi^e_{j,t}) + \delta_2(T_i \times \pi^e_{j,t}) + \delta_3(T_i \times T_j \times \pi^e_{j,t}) \\
&+ \tilde{\pi}^2_t \times \left[\beta^\pi_0 g_{j,t} + \beta^\pi_1(T_j \times g_{j,t}) + \beta^\pi_2(T_i \times g_{j,t}) + \beta^\pi_3(T_i \times T_j \times g_{j,t})\right] \\
&+ \tilde{\pi}^2_t \times \left[\delta^\pi_0 \pi^e_{j,t} + \delta^\pi_1(T_j \times \pi^e_{j,t}) + \delta^\pi_2(T_i \times \pi^e_{j,t}) + \delta^\pi_3(T_i \times T_j \times \pi^e_{j,t})\right] + \varepsilon_{ij,t}. \qquad (80)
\end{aligned}
$$

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 6 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 7 cleanly separates them.

Trade value and buyer beliefs. Start from CES demand (14) and write bilateral trade value as

$$V_{k\ell,t} \;=\; p_{k\ell,t} \cdot x_{k\ell,t}(p_{k\ell,t}) \;=\; p_{k\ell,t}^{1-\varepsilon_i}\; \mathbb{E}[P_{j,t} \mid \mathcal{I}_{k,t}]^{\varepsilon_i - 1}\; a_{ij}\, S_{i,t}, \tag{81}$$

where k denotes the buyer (in sector i), ℓ the supplier (in sector j), $P_{j,t}$ the supplier-sector cost shifter, and $\mathcal{I}_{k,t}$ the buyer's information set. Taking logs,

$$\log V_{k\ell,t} \;=\; (1-\varepsilon_i)\, \log p_{k\ell,t} \;+\; (\varepsilon_i - 1)\, \log \mathbb{E}[P_{j,t} \mid \mathcal{I}_{k,t}] \;+\; \log(a_{ij}\, S_{i,t}). \tag{82}$$

Bilateral trade value depends on the supplier's price (which reduces demand through substitution) and the buyer's conditional expectation of the outside option $P_{j,t}$ (which sustains demand when the buyer believes costs are high industry-wide). Both the price and the belief channels operate through $\log \mathbb{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$: any signal that shifts the buyer's expectation of $P_{j,t}$ affects bilateral trade.

Two channels of information transmission. The buyer's conditional expectation $\mathbb{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$ 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 (15). The second is the supplier's communicated belief $\hat{z}_{j\ell,t}$, 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 (17)), 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 $\omega_{ij,t}$:

$$\frac{\partial \log \mathbb{E}[P_{j,t} \mid \mathcal{I}_{k,t}]}{\partial \log p_{k\ell,t}} \;=\; \omega_{ij,t}.$$

Substituting into (82), the trade-value elasticity with respect to the bilateral price is

$$\frac{\partial \log V_{k\ell,t}}{\partial \log p_{k\ell,t}} \;=\; -(\varepsilon_i - 1)(1 - \omega_{ij,t}). \tag{83}$$

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 $P_{j,t}$, dampening substitution. The treatment differential in the trade-value elasticity is

$$\beta_1^p \;=\; (\varepsilon_i - 1)\, \Delta\omega_{ij} \;>\; 0, \tag{84}$$

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 $\hat{z}_{j\ell,t}$ 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 $\hat{z}_{j\ell,t}$.

The trade-value response to the supplier's belief is

$$\frac{\partial \log V_{k\ell,t}}{\partial \hat{z}_{j\ell,t}} \;=\; (\varepsilon_i - 1)\,\psi_{ij,t} \;-\; (\varepsilon_i - 1)\,\frac{\partial \log p_{k\ell,t}}{\partial \hat{z}_{j\ell,t}}, \tag{85}$$

where the first term is the demand-sustaining effect (higher $\hat{z}_{j\ell,t}$ raises $\mathbf{E}[P_{j,t} \mid \mathcal{I}_{k,t}]$, 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 $\hat{z}_{j\ell,t}$ from (13), reducing demand). The treatment differential is

$$\beta^e_1 \;=\; (\varepsilon_i - 1)\Big(\psi^{TR}_{ij,t} - \psi^{UT}_{ij,t}\Big) \;>\; 0, \tag{86}$$

which is positive because treatment reduces $\sigma^2_{\nu,j}$, making $\hat{z}_{j\ell,t}$ more precise and raising the weight $\psi_{ij,t}$ the buyer places on it.

Equations (84) and (86) 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 3), 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 5.5 includes both channels simultaneously:

$$\begin{aligned}\log V_{k\ell,t} = \alpha_k + \delta_{s(j),t} + \alpha_{T_j,t} + \beta^p_0\, g^p_{j,t} + \beta^p_1\,(T_j \times g^p_{j,t}) \\ + \beta^e_0\, \pi^e_{j,t} + \beta^e_1\,(T_j \times \pi^e_{j,t}) + \varepsilon_{k\ell,t},\end{aligned} \tag{87}$$

where $\delta_{s(j),t}$ are supplier-industry×quarter fixed effects and the remaining notation follows equation (25). 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 $\sigma^2_z/\sigma^2_\zeta$ 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 7) 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 7. 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 7, 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 7, 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 7, 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 $R_{ij} = 1/(1 - \beta\rho_{ij})$, 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.

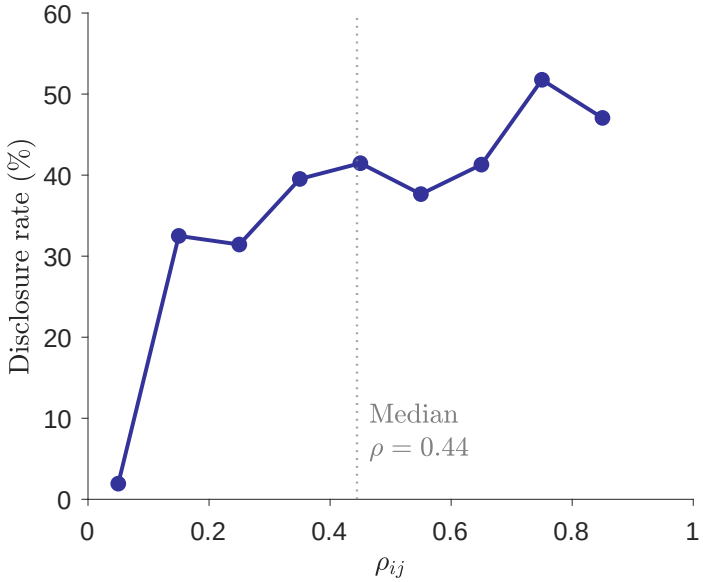


Figure R.1: Equilibrium disclosure rate by relationship persistence

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 $\hat{\sigma}^2_{z,j}$ (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 (ψ, λ_ζ), calibrated by minimum distance. The disclosure cost is a uniform setup cost ψ 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 (26). The coefficient $\hat{\phi}_j$ estimates the Kalman gain $\phi_j = \sigma^2_z/(\sigma^2_z + \sigma^2_\nu)$, from which $\hat{\sigma}^2_{\nu,j} = \hat{\sigma}^2_{z,j}(1 - \hat{\phi}_j)/\hat{\phi}_j$. 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 ($\hat{\phi} = 0.48$). 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 $\sigma^2_{z,j}$ is estimated from the time-series variance of industry-level output price inflation.

Disclosure setup cost. ψ is the per-link cost of establishing a verifiable disclosure protocol. It is identified by the upstream belief premium (moment m1): a higher ψ 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.

$$\hat{\beta}_1^{\text{up}} = \text{Cov}(T_j,\, \pi^e_{i,t})/\text{Var}(T_j), \tag{88}$$

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.

$$\hat{\beta}_2 = \text{Cov}(T_j \cdot g^p_{j,t},\, \pi^e_{i,t})/\text{Var}(T_j \cdot g^p_{j,t}), \tag{89}$$

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 $\omega_{p,j} = \sigma^2_{z,j}/(\sigma^2_{z,j} + \lambda_\zeta\hat{\sigma}^2_{\zeta,j})$.

S.3 Objective function

The calibration minimizes a diagonal GMM criterion over the two targeted moments:

$$\mathcal{L}(\theta) = \sum_{k=1}^{2} \left(\frac{m_k(\theta) - \hat{m}_k}{\hat{s}_k} \right)^2, \tag{90}$$

where $m_k(\theta)$ is the model-implied analog of moment k at parameters $\theta = (\psi, \lambda_\zeta)$, $\hat{m}_k$ 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 $\log_{10}(\psi) \in [-12, -4]$ 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 $\mathcal{W}_1$.

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 ψ 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 $\psi/(R_{ij} \cdot W(\omega_{p,j}) \cdot \Gamma_{ij})$—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 ψ inherits these units and is therefore small in absolute terms by construction.

Cost-to-surplus ratios. The economically meaningful quantity is the ratio $\psi/(R_{ij} \cdot W(\omega_{p,j}) \cdot \Gamma_{ij})$, 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 ψ is a pure unit artefact of GDP normalisation. The economically meaningful object is the ratio $\psi/(R_{ij} \cdot W(\omega_j) \cdot \Gamma_{ij})$, 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 6.

T Fixed-effect robustness

This appendix examines the sensitivity of the baseline learning results (Table 3) to alternative fixed-effect structures. Table L.18 (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 (price 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 (price 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.