

INTERNATIONAL MONETARY FUND

Production Networks and the (Asymmetric) Transmission of Monetary Policy

Francesco Grigoli

WP/26/127

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2026
JUN**



WORKING PAPER

IMF Working Paper
Research Department

Production Networks and the (Asymmetric) Transmission of Monetary Policy

Prepared by Francesco Grigoli*

Authorized for distribution by Antonio Spilimbergo
June 2026

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: I study how the production network shapes monetary policy transmission to prices. Using U.S. data, I show that industries farther upstream from final demand exhibit larger cumulative price responses to monetary shocks, while downstream industries absorb shocks through output. A calibrated multi-sector New Keynesian model rationalizes these patterns: upstream sectors, which sell predominantly to other firms, reprice more frequently and therefore exhibit less price rigidity. A counterfactual decomposition of the price response shows that this heterogeneity in price rigidity—rather than cost-cascade propagation through input-output linkages—is the primary driver of the cross-sectional responses. The upstreamness differential is strongly asymmetric, large following expansionary shocks but nearly absent following contractionary ones, consistent with asymmetric price rigidity compounding across production stages. Together, these findings suggest that monetary policy's potency depends on the production network's architecture.

RECOMMENDED CITATION: Grigoli, Francesco “Production Networks and the (Asymmetric) Transmission of Monetary Policy” IMF Working Paper No. 26/127, June 2026

JEL Classification Numbers:	E31, E52, L14, C67
Keywords:	production networks; price rigidity; monetary policy transmission; input-output linkages; asymmetric price adjustment
Author's E-Mail Address:	francesco.grigoli@georgetown.edu

* “The author would like to thank Emiliano Luttini and colleagues at Georgetown University for discussions and comments on earlier drafts. This paper was completed while the author was with the IMF's Research Department.

1 Introduction

Every dollar of GDP passes through an average of 2.1 intermediate transactions before reaching a final consumer.¹ This production network structure is a powerful propagation mechanism for real shocks, as idiosyncratic sectoral disturbances can cascade through input-output (IO) linkages and generate aggregate fluctuations ([Acemoglu *et al.*, 2012](#); [Baqaee and Farhi, 2019](#); [Carvalho and Tahbaz-Salehi, 2019](#)). Whether production networks also shape the transmission of nominal shocks is far less understood. When the Federal Reserve changes the federal funds rate, the impulse must travel through the web of buyer-seller relationships, encountering a distinct pricing decision—and a distinct degree of nominal rigidity—at each node. Yet the workhorse New Keynesian model collapses this structure into a single representative firm facing a single Calvo friction; and, despite the evidence that prices adjust more readily upward than downward ([Peltzman, 2000](#); [Nakamura and Steinsson, 2008](#)), it assumes symmetric price adjustment for rate hikes and cuts. This paper asks whether the architecture of the production network shapes the transmission of monetary policy to prices, and whether it does so differently for expansionary and contractionary shocks.

I develop a multi-sector New Keynesian model in which industries are connected through an IO network and each faces Calvo-style nominal rigidity. The key theoretical result is that aggregate monetary non-neutrality depends on a network-weighted rigidity index rather than the simple average of sectoral price stickiness. As a result, two economies with identical average price rigidity but different production structures will exhibit different degrees of monetary non-neutrality.² The model delivers a sharp cross-sectional prediction: more upstream industries—those farther from final demand, as measured by the upstreamness index of [Antràs *et al.* \(2012\)](#)—should exhibit larger cumulative price responses to monetary shocks. This is because upstream producers sell predominantly to other firms rather than to consumers, and business-to-business prices are renegotiated more frequently than consumer-facing prices ([Nakamura and Steinsson, 2008](#)). Network position thus generates a systematic correlation between upstreamness and intrinsic price rigidity, with more upstream sectors having lower Calvo parameters.³

I test that more upstream industries exhibit larger cumulative price responses to monetary shocks using a panel of monthly producer price indices for 271 U.S. industries at the six-digit NAICS level and the orthogonalized high-frequency monetary policy surprises of [Bauer and Swanson \(2023\)](#). Local projections of cumulative industry price changes on the interaction of the mone-

¹Computed from the Bureau of Economic Analysis detailed input-output tables as the sales-weighted average upstreamness across all industries, following [Antràs *et al.* \(2012\)](#).

²The model shares the multi-sector structure of [Pasten *et al.* \(2020\)](#) and [Rubbo \(2023\)](#), but shifts the focus from the aggregate statistic to the cross-sectional anatomy of monetary transmission—asking which industries adjust and through what channel.

³Consider a steel mill that sells to an auto parts manufacturer, who in turn sells to an automobile assembler. A monetary contraction reduces demand along the entire chain, but the steel mill’s price adjustment—being three steps from the final consumer—must pass through two intermediate pricing decisions before reaching retail prices. If either the parts manufacturer or the assembler has sticky prices, the steel mill’s price reduction is partially absorbed at each stage.

tary shocks with upstreamness reveal that the differential price response per unit of upstreamness is statistically and economically significant starting from one year after the shock up to over two years after the shock, peaking one year and a half after the shock. The empirical results imply that an industry at the 75th percentile of upstreamness exhibits a peak cumulative price response 2.8 percentage points larger than one at the 25th percentile, in response to a one percentage point monetary policy shock. Feeding the Bureau of Economic Analysis (BEA) IO matrix into the model with an externally measured rigidity differential—calibrated from industry-level PPI price adjustment frequencies—generates the correct qualitative pattern of interaction coefficients across horizons and captures a substantial share of the quantitative magnitude, providing a credible link between the theoretical mechanism and the reduced-form evidence.

A counterfactual decomposition of the calibrated model confirms that this network-determined rigidity heterogeneity—which I call the *flexibility channel*—is the primary driver of the cross-sectional pattern. Instead, the cost propagation, or *cost-cascade channel*, usually emphasized in multi-sector models—through which upstream price adjustments propagate to downstream marginal costs—plays only a secondary, and partially offsetting, role. The intuition is that, holding rigidity fixed, a sector’s price tracks the flexible wage in proportion to its labor share, so cost linkages would reinforce the upstreamness pattern only if upstream sectors were labor-intensive. In U.S. data they are the opposite—intermediate-input intensive, with marginal costs tied to other, themselves sticky, prices rather than to wages. The cost-cascade channel therefore tilts the larger price responses toward downstream sectors, working against the upstream flexibility advantage rather than amplifying it. This decomposition clarifies that the production network’s primary role is to determine which sectors are flexible, rather than to propagate price adjustments across stages.

I highlight a stark asymmetry in the cross-sectional pattern. Extending the model to allow for asymmetric price rigidity—the empirical regularity that prices adjust more readily upward than downward—I show that the upstreamness differential in price responses is substantially larger following expansionary shocks than contractionary ones. The data are consistent with this prediction. In response to a rate cut, the peak interaction coefficient is more than twice the pooled peak; cumulating over the one-to-two-year window where the responses diverge most, the upstreamness differential is large and negative following rate cuts but near zero following rate hikes, providing evidence of downward nominal rigidity. This corroborates the ‘rockets and feathers’ phenomenon long documented at the aggregate level, indicating that inflation is easier to generate than deflation. For central banks, the implication is that rate cuts generate larger cross-sectional price dispersion than rate hikes—a structural asymmetry in the transmission mechanism that is invisible in representative-firm models.

These price dynamics in response to monetary policy shocks have implications for the real economy. The model highlights that industries whose prices adjust more flexibly absorb monetary shocks primarily through price changes, preserving real output; industries with sticky prices instead absorb shocks through quantity changes. The empirical results match these predictions. Upstream sectors, whose prices are more flexible, exhibit less negative real effects than down-

stream sectors. Examining contractionary and expansionary shocks separately reveals patterns mirroring the price-side asymmetry. That is, I do not find statistically significant real effects in response to expansionary shocks. However, contractionary shocks lead upstream industries to cut prices more aggressively and preserve real output, and downstream industries to absorb the shock through quantity rationing. In sum, price flexibility substitutes for quantity adjustment, generating cross-sectional dispersion in both prices and output that is concentrated in contractionary episodes, consistent with the asymmetric price rigidity mechanism in the model.

Although the decomposition shows that the cost-cascade channel is not the primary driver of the upstreamness differential, costs do propagate through the network, and this propagation is detectable in the data. Three auxiliary tests support the view that buyer-seller price linkages play a role in transmitting price adjustments across sectors, even if they do not generate the cross-sectional pattern on their own. First, controlling for input-cost-weighted supplier prices absorbs most of the upstreamness interaction, consistent with the effect operating partly through cost linkages rather than omitted industry characteristics. Second, the interaction between upstreamness and the monetary shock is significantly stronger for industries with higher intermediate input shares. Third, an industry’s suppliers’ network positions predict its own price response.

These findings survive extensive robustness checks. Alternative network measures (eigenvector centrality, intermediate input share) and monetary policy shocks ([Jarociński and Karadi, 2020](#)) yield significant interaction coefficients of comparable magnitude. Similarly, controlling for industry characteristics potentially correlated with industry upstreamness delivers consistent results. The results are also robust to constraining the sample to the manufacturing sector or to removing one industry at a time. Finally, permutation-based placebo tests indicate that upstreamness—not an arbitrary industry ranking—generates the differential price response.

The results of this paper have implications for monetary policy analysis. If the cross-sectional distribution of price rigidity is shaped by the production network, then structural changes in the network can alter the potency of monetary policy even if average price rigidity remains unchanged. Any policy or structural shift that alters the IO matrix—trade agreements, vertical integration, platform-mediated disintermediation—changes the composition of B2B versus B2C activity across the economy, reshaping the distribution of rigidity across network positions. The secular increase in the services share of GDP shifts economic activity toward sectors that are both more downstream and more price-rigid, increasing the wedge between upstream and downstream rigidity and potentially strengthening the asymmetry documented in this paper.

1.1 Related literature and contributions

This paper contributes to three strands of the literature. The first is the literature on production networks and nominal rigidities. A large body of work has established that IO linkages shape the propagation of real shocks ([Long and Plosser, 1983](#); [Acemoglu *et al.*, 2012](#); [Gabaix, 2011](#); [Barrot and Sauvagnat, 2016](#); [Carvalho *et al.*, 2021](#); [Baqaee and Farhi, 2019](#)).⁴ A separate literature has studied

⁴[Carvalho \(2014\)](#) and [Carvalho and Tahbaz-Salehi \(2019\)](#) provide surveys.

how heterogeneity in nominal price rigidity affects aggregate monetary non-neutrality (Carvalho, 2006; Nakamura and Steinsson, 2010; Kehoe and Midrigan, 2015).⁵

The intersection of these two strands of literature—how nominal rigidities interact with production network structure—has received less attention. The theoretical foundation was laid by Basu (1995), who showed that when firms use other firms’ outputs as intermediate inputs, each firm’s marginal cost depends on its suppliers’ prices, so that upstream price rigidity propagates downstream and amplifies aggregate monetary non-neutrality. Building on this insight, Nakamura and Steinsson (2010) develop a multi-sector menu cost model with intermediate inputs calibrated to 14 U.S. sectors, showing that the combination of heterogeneous price stickiness and roundabout production substantially amplifies monetary non-neutrality. Pasten *et al.* (2020) calibrate a multi-sector Calvo model with sector-specific adjustment probabilities and find that heterogeneity in price-adjustment frequency is the dominant driver of monetary non-neutrality while IO linkages contribute only marginally—the aggregate counterpart of the cross-sectional decomposition I develop below. Relatedly, Pastén *et al.* (2024) show that heterogeneity in sectoral price stickiness amplifies the aggregate effects of idiosyncratic productivity shocks and reshapes which sectors matter for aggregate fluctuations. Rubbo (2023) demonstrates that IO linkages propagate rigidity across stages, flattening the aggregate Phillips curve. More recently, Afrouzi and Bhattarai (2024) derive closed-form sufficient statistics for inflation and GDP dynamics in this class of models, showing that all dynamics are governed by a duration-adjusted Leontief matrix and that, as in my decomposition, an upstream sector’s contribution depends on whether it is flexible or sticky, not on its network position alone. My theoretical contribution relative to this literature, which uniformly assumes symmetric price adjustment, is to extend the propagation matrix to sign-dependent rigidity. That is, asymmetric rigidity splits the duration-adjusted Leontief into distinct expansionary and contractionary matrices, whose gap generates a network-amplified asymmetry that is absent from all symmetric frameworks.

All these papers focus on aggregate statistics—the level of GDP volatility, the slope of the Phillips curve. Ghassibe (2023) is the first to provide empirical evidence on network amplification of monetary policy, estimating that at least 30% of the effect of monetary shocks on U.S. aggregate final consumption comes from amplification through IO linkages. His analysis uses sectoral consumption data for up to 161 BEA sectors and focuses on the aggregate contribution of networks. I shift the lens from the aggregate to the cross-section, asking not how much total non-neutrality production networks generate but which industries bear the largest price and output responses and through what mechanism. I provide direct empirical evidence for the cross-sectional anatomy of monetary transmission through production chains and document a novel asymmetry that is absent from all three frameworks.

The second strand of literature this paper relates to is the one on asymmetric price adjust-

⁵Note the distinction between Carlos Carvalho, who studies heterogeneous price rigidity and monetary non-neutrality (Carvalho, 2006), and Vasco Carvalho, who studies production networks (Carvalho, 2014; Carvalho and Tahbaz-Salehi, 2019). Carvalho and Schwartzman (2015)—by Vasco Carvalho and Schwartzman—bridges both literatures, studying how selection effects in time-dependent pricing models affect monetary non-neutrality.

ment and production networks. Micro-level evidence has long documented that prices rise faster than they fall (Peltzman, 2000; Nakamura and Steinsson, 2008). In a representative-firm model, such asymmetric rigidity merely shifts the level of price adjustment. I show that in a production network, asymmetric rigidity shifts the differential across upstream and downstream sectors: the asymmetry compounds through $K - 1$ intermediate pricing decisions along a chain of length K , so the direction of the shock matters multiplicatively. The data are in line with this prediction. This result connects to the broader literature on nonlinearities in monetary transmission (Tenreyro and Thwaites, 2016) and provides a micro-founded, cross-sectional mechanism for the aggregate observation that inflation is easier to generate than deflation.

Finally, this paper also speaks to the empirical evidence on networks and monetary policy. Ozdagli and Weber (2025) document that network-central firms exhibit stronger stock-price responses to monetary surprises, and Ghassibe (2023) provides consumption-side evidence that network amplification accounts for a substantial share of aggregate monetary non-neutrality. I focus on goods prices, providing a more direct test of the New Keynesian pricing mechanism.⁶ Moreover, I provide suggestive evidence that costs propagate through buyer-seller price linkages, confirming that the IO network transmits price signals across production stages even though the cross-sectional differential is primarily driven by the flexibility channel. Closest in spirit on the empirical side is Minton and Wheaton (2023), who use sectoral PPI data and the BEA IO tables to show that commodity-price shocks propagate downstream with delays that depend on sectoral price flexibility. My focus differs in the shock (identified monetary surprises rather than commodity prices), the object (the cross-sectional response by upstreamness), and the asymmetry between expansionary and contractionary shocks. Relatedly, the paper connects to the empirical literature on substantial heterogeneity in price adjustment frequencies (Bils and Klenow, 2004; Nakamura and Steinsson, 2008). I show that this heterogeneity is systematically—though imperfectly—related to production chain position: upstream sectors tend to be more flexible because they sell predominantly to other firms, but the relationship is far from deterministic, as some upstream sectors (e.g., semiconductor manufacturing) are among the stickiest while some downstream sectors (e.g., food) are among the most flexible (Pasten *et al.*, 2020).

Outline Section 2 develops the theoretical model. Section 3 describes the data. Section 4 presents the empirical strategy. Section 5 reports the main results of the paper and mechanism tests. Section 6 presents the results of the robustness tests. Section 7 concludes.

2 Model

I develop a multi-sector New Keynesian model with IO linkages, building on the canonical framework of Woodford (2003) and Galí (2015). The model is deliberately stylized to isolate two mech-

⁶Stock prices reflect both price and quantity adjustments. PPI data isolate the price margin, mapping directly to sectoral Phillips curves. See also La’O and Tahbaz-Salehi (2022), who characterize optimal monetary policy in production networks.

anisms: (i) network-determined heterogeneity in nominal rigidity, whereby upstream sectors—selling predominantly to other firms—have more flexible prices, which generates a cross-sectional pattern of price responses with respect to upstreamness, and (ii) asymmetric nominal rigidity interacting with this network-determined rigidity differential, which causes the cross-sectional pattern to differ sharply between inflationary and deflationary shocks. The second mechanism is the model’s most distinctive prediction and cannot arise in representative-firm models.⁷

2.1 Environment

Households A representative household maximizes lifetime utility

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma}}{1-\sigma} - \frac{L_t^{1+\varphi}}{1+\varphi} \right], \quad (1)$$

where C_t is a composite consumption good, L_t is labor supply, $\sigma > 0$ is the coefficient of relative risk aversion, $\varphi > 0$ is the inverse Frisch elasticity, and $\beta \in (0, 1)$ is the discount factor. The composite consumption good aggregates output from N sectors:

$$C_t = \left(\sum_{i=1}^N \omega_i^{1/\eta} c_{it}^{(\eta-1)/\eta} \right)^{\eta/(\eta-1)}, \quad (2)$$

where c_{it} is consumption of sector- i output, ω_i is the expenditure weight on sector i in final demand, and $\eta > 0$ is the elasticity of substitution across sectors.

Firms Each sector $i \in \{1, \dots, N\}$ contains a continuum of monopolistically competitive firms indexed by $j \in [0, 1]$. Firm j in sector i produces output using labor and intermediate inputs from all N sectors:

$$y_{ijt} = z_{it} \ell_{ijt}^{\alpha_i} \prod_{k=1}^N m_{ijk,t}^{\gamma_{ik}}, \quad (3)$$

where z_{it} is a sector-wide productivity shifter, ℓ_{ijt} is labor input, $m_{ijk,t}$ is the quantity of sector- k output used as an intermediate input, and the production parameters satisfy $\alpha_i + \sum_{k=1}^N \gamma_{ik} = 1$ for all i .

The matrix $\Gamma = [\gamma_{ik}]$ is the IO matrix, where γ_{ik} is sector i ’s cost share of inputs from sector k . This matrix encodes the production network.

Price setting Firms face Calvo-style nominal rigidity (Calvo, 1983).⁸ In each period, a fraction $1 - \theta_i$ of firms in sector i can reset their price; the remaining fraction θ_i must keep their price un-

⁷Appendix A-I reports the proofs of the model propositions.

⁸The Calvo specification is chosen for analytical tractability. State-dependent pricing models (e.g., Golosov and Lucas, 2007; Midrigan, 2011) would introduce additional nonlinearities but would not change the qualitative prediction that network position matters for the cross-section of price responses.

changed. The parameter θ_i varies across sectors, capturing the well-documented heterogeneity in price adjustment frequencies (Nakamura and Steinsson, 2008; Bils and Klenow, 2004). I initially present the symmetric Calvo case for transparency, then extend the model to allow the reset probability to depend on the direction of the desired price change. This extension is central to the paper's main result. Since asymmetric rigidity interacts with the systematic variation in price rigidity across network positions, the cross-sectional effect operates more strongly for inflationary than for deflationary shocks.

When a firm in sector i can reset its price, it chooses p_{ijt}^* to maximize the expected discounted sum of profits:

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta \theta_i)^s \left[p_{ijt}^* y_{ij,t+s} - \text{TC}_{ij,t+s}(y_{ij,t+s}) \right], \quad (4)$$

where $\text{TC}_{ij,t+s}(\cdot)$ is the total cost function at date $t+s$, taking input prices as given.

Monetary authority The central bank sets the nominal interest rate according to a Taylor rule with monetary policy shocks:⁹

$$i_t = \rho_r i_{t-1} + (1 - \rho_r) [\phi_\pi \pi_t^{\text{agg}} + \phi_y \hat{y}_t] + \varepsilon_t^m, \quad (5)$$

where i_t is the nominal interest rate, π_t^{agg} is aggregate inflation, the Domar-weighted average of sectoral inflation, $\hat{y}_t$ is the output gap, and ε_t^m is an exogenous monetary policy shock.

2.2 Log-linearized equilibrium

Log-linearizing around the zero-inflation steady state yields a system of sectoral New Keynesian Phillips curves linked through the IO structure.

Proposition 1 (Sectoral Phillips Curves). *The log-linearized inflation rate in sector i satisfies*

$$\pi_{it} = \beta \mathbb{E}_t \pi_{i,t+1} + \lambda_i \left[\alpha_i \hat{w}_t + \sum_{k=1}^N \gamma_{ik} \hat{p}_{kt} - \hat{p}_{it} \right], \quad (6)$$

where $\hat{w}_t$ is the log deviation of the nominal wage from steady state, $\hat{p}_{it}$ is the log deviation of sector i 's price index, and $\lambda_i \equiv \frac{(1-\theta_i)(1-\beta\theta_i)}{\theta_i}$ is the slope of sector i 's Phillips curve—a decreasing function of the Calvo parameter θ_i .

The result follows from standard Calvo aggregation applied to each sector. Firms that can reset choose p_{ijt}^* to equate the present-value-weighted average of expected markups to the desired markup. Log-linearizing the optimal reset price and aggregating across resetting and non-resetting firms using the CES price index yields Equation 6. The Cobb-Douglas production func-

⁹The Taylor rule is a modeling device to generate exogenous monetary policy variation. The empirical analysis uses externally identified shocks and does not rely on the specific functional form of the policy rule.

tion in Equation 3 delivers the log-linear marginal cost expression in brackets, which corresponds to the log deviation of sector i 's real marginal cost from its flexible-price level. It depends on the wage (through the labor share α_i) and on the prices of all intermediate inputs (through the IO weights γ_{ik}). This is the channel through which the production network enters the inflation dynamics, as sector i 's marginal cost depends on the prices set by its suppliers, which in turn depend on their suppliers' prices, and so on through the chain.

Matrix representation Stacking the sectoral Phillips curves, define the vectors $\pi_t = (\pi_{1t}, \dots, \pi_{Nt})'$ and $\hat{p}_t = (\hat{p}_{1t}, \dots, \hat{p}_{Nt})'$. Let $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_N)$ and $\mathcal{A} = \text{diag}(\alpha_1, \dots, \alpha_N)$. Then:

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \Lambda [\mathcal{A} \hat{w}_t \mathbf{1} + (\Gamma - \mathbf{I}) \hat{p}_t]. \quad (7)$$

This system reveals how the network enters the inflation dynamics. The IO matrix Γ creates cross-equation linkages among the sectoral Phillips curves, so that even with identical Calvo parameters across sectors, the network structure would generate heterogeneous inflation dynamics. This observation connects directly to the “network Phillips curve” of Rubbo (2023), who shows that in the aggregate the trade-off between inflation and output depends on the entire spectrum of the IO matrix, not merely on the average rigidity parameter.¹⁰

2.3 Aggregate non-neutrality and network structure

The central theoretical result characterizes how aggregate monetary non-neutrality depends on the production network.

Definition 1 (Domar weights). *The Domar weight of sector i , denoted ξ_i , is the ratio of sector i 's total sales (to both final consumers and other firms) to aggregate GDP. In the steady state, the vector of Domar weights satisfies $\xi = (\mathbf{I} - \Gamma')^{-1} \omega$, where ω is the vector of final demand shares.*¹¹

Definition 2 (Upstreamness). *The upstreamness of sector i , following Antràs et al. (2012), is the expected number of production stages between sector i and final demand:*

$$U_i = \mathbf{e}_i' (\mathbf{I} - \mathbf{A})^{-1} \mathbf{1}, \quad (8)$$

where $\mathbf{e}_i$ is the i -th unit vector and $\mathbf{A}$ is the output-allocation matrix, with $a_{ik} = Z_{ik}/Y_i$ denoting the share of sector i 's output sold to sector k (as intermediate input or final demand). A sector with $U_i = 1$ sells entirely to final consumers; higher values indicate more upstream positions.¹²

¹⁰My framework shares this algebraic structure but emphasizes the cross-sectional predictions rather than the aggregate trade-off. Appendix A-II derives the network Phillips curve.

¹¹Domar weights can exceed unity because gross output includes intermediate sales. In U.S. data, the sum of Domar weights across all industries is approximately 1.8, reflecting the substantial role of intermediate transactions in total economic activity.

¹²The output-allocation matrix $\mathbf{A}$ is distinct from the cost-share matrix Γ used in the production function. While γ_{ik} captures sector i 's expenditure share on inputs from k (a demand-side object), a_{ik} captures the share of sector i 's sales

Proposition 2 (Network-adjusted rigidity). *The aggregate price level response to a monetary shock ε_t^m at horizon h is*

$$\frac{\partial \hat{P}_{t+h}}{\partial \varepsilon_t^m} = \sum_{i=1}^N \xi_i \frac{\partial \hat{p}_{i,t+h}}{\partial \varepsilon_t^m}, \quad (9)$$

where $\hat{P}_t = \sum_i \xi_i \hat{p}_{it}$ is the Domar-weighted aggregate price level. The real output effect of the monetary shock is, to a first-order approximation,

$$\hat{Y}_{t+h} \approx \hat{D}_{t+h} - \hat{P}_{t+h}, \quad (10)$$

so that output non-neutrality is larger when $\hat{P}_{t+h}$ responds slowly to the monetary shock—that is, when the network-weighted rigidity is high.

The aggregate price index $\hat{P}_t = \sum_i \xi_i \hat{p}_{it}$ is the Domar-weighted gross-output price index.¹³ Equation 9 follows by linearity. The output approximation uses the aggregate demand identity in log-deviations: $\hat{Y}_{t+h} \equiv \hat{D}_{t+h} - \hat{P}_{t+h}$, where $\hat{D}_{t+h}$ denotes the nominal aggregate demand impulse from the monetary shock (i.e., the path of nominal spending implied by the Taylor rule and Euler equation, not a money stock).

Corollary 1 (Centrality-rigidity alignment). *The aggregate price level is slowest to adjust—and monetary non-neutrality is greatest—when the correlation between sectoral rigidity θ_i and Domar weight ξ_i is positive. That is, monetary policy has the largest real effects when the stickiest-price sectors are also the most central in the production network.*

This result generalizes the standard Calvo model in that in a one-sector economy, the single θ parameter determines non-neutrality. In a network economy, what matters is not $\bar{\theta}$ but the Domar-weighted average $\tilde{\theta} \equiv \sum_i \xi_i \theta_i$. The insight that the alignment between Domar weights and sectoral rigidity governs aggregate non-neutrality is closely related to the decomposition in [Pasten et al. \(2020\)](#), who separate the contributions of IO linkages and Calvo heterogeneity to aggregate monetary non-neutrality.

Whether the alignment condition in [Corollary 1](#) holds empirically is an important quantitative question. In U.S. data, the evidence suggests that retail-facing sectors (food, apparel) tend to have both higher Domar weights and stickier consumer prices, while upstream commodity-producing sectors (chemicals, primary metals) have lower Domar weights and more flexible prices. [Nakamura and Steinsson \(2008\)](#) and [Bils and Klenow \(2004\)](#) document substantial heterogeneity in the

going to k (a supply-side object). Upstreamness measures forward linkages—how many stages of production a sector’s output must pass through before reaching final consumers—and therefore requires the sales-share matrix $\mathbf{A}$, not the cost-share matrix Γ .

¹³Because the sum of Domar weights exceeds unity ($\sum_i \xi_i \approx 1.8$ in U.S. data), $\hat{P}_t$ is a gross-output-weighted price index that counts intermediate transactions. This is the natural price aggregate for analyzing how sectoral price stickiness translates into aggregate non-neutrality, following the approach in [Baqaee and Farhi \(2020\)](#). The GDP deflator—the final-demand-weighted price index $\sum_i \omega_i \hat{p}_{it}$ —would assign zero weight to intermediate stages and miss the network amplification channel.

frequency of price changes, and cross-referencing their estimates with the BEA IO tables suggests a positive but moderate correlation between θ_i and ξ_i . The model therefore predicts that the U.S. production network amplifies monetary non-neutrality relative to a counterfactual economy with the same average rigidity but no network structure.

Cross-sectional predictions The model generates testable predictions about the cross-section of sectoral price responses. These predictions rest on a maintained assumption, motivated by the empirical regularity that inter-firm (i.e., B2B) prices are renegotiated more frequently than consumer-facing prices (Nakamura and Steinsson, 2008):

Assumption 1. *More upstream sectors have weakly more flexible prices: θ_i is weakly decreasing in U_i , so that $\lambda_i = (1 - \theta_i)(1 - \beta\theta_i)/\theta_i$ is weakly increasing in U_i .*

Equivalently, the price-adjustment probability $r_i \equiv \lambda_i/(1 + \lambda_i)$ is weakly increasing in upstreamness. The assumption is motivated by the observation that upstream sectors sell predominantly to other firms, and B2B prices are more flexible. An important caveat is that the share of sales to other firms is not identical to upstreamness. That is, two sectors can both sell entirely to firms yet differ in upstreamness if they sit at different positions in the production chain. The assumption therefore requires that flexibility is monotone in the upstreamness index U_i —not only in the B2B sales share—which is an empirical restriction.

Proposition 3 (Upstreamness and price dynamics). *Under Assumption 1 and in the static case ($\beta = 0$), following a contractionary monetary shock ($\epsilon_t^m > 0$):*

(i) **General network.** *The impact price response of sector i decomposes as*

$$\frac{\partial \hat{p}_{it}}{\partial \epsilon_t^m} = r_i \times \widehat{mc}_i, \quad (11)$$

where $r_i = \lambda_i/(1 + \lambda_i)$ is sector i 's price-adjustment probability and $\widehat{mc}_i > 0$ is its marginal cost response. The flexibility term r_i is increasing in U_i by Assumption 1, which provides a force toward larger price responses for more upstream sectors. Whether the total response $r_i \times \widehat{mc}_i$ is monotone in U_i also depends on how the marginal cost response $\widehat{mc}_i$ varies with network position, which is a quantitative question I resolve in the calibration.

(ii) **Chain economy: prices.** *In a chain of K sectors with uniform input share γ and uniform r , the price response of sector k is*

$$x_k = \frac{r \left[(1 - \gamma) + \gamma(1 - r) \underbrace{(r\gamma)^{k-1}}_{\phi} \right]}{1 - r\gamma}, \quad (12)$$

which is strictly decreasing in k (the sector's position toward final demand). The upstream-downstream price differential $x_1 - x_K$ is

$$D_K = \frac{r\gamma(1-r) [1 - \phi^{K-1}]}{1 - \phi}, \quad (13)$$

where $\phi = r\gamma$ is the cascade factor. The differential is increasing in K and converges to $r\gamma(1-r)/(1-\phi)$ as $K \rightarrow \infty$.

- (iii) **Chain economy: output.** The real output response of sector k satisfies $\hat{y}_{kt} \approx \hat{d}_{kt} - \hat{p}_{kt}$, where $\hat{d}_{kt}$ is the demand impulse. More downstream sectors (higher k) have smaller $|\hat{p}_{kt}|$, so their output gap remains open longer and the peak real effect occurs at a later horizon.

Part (i) holds for an arbitrary IO matrix Γ . The factorization in Equation 11 follows from the static Phillips curve, so that sector i 's price adjusts by a fraction r_i of its marginal cost deviation. The marginal cost $\widehat{mc}_i = \alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt}$ is strictly positive because both the wage and all input prices respond in the same direction as the monetary shock, and $\alpha_i > 0$. Two conceptually distinct channels shape the cross-sectional pattern. The flexibility channel operates through r_i : upstream sectors have steeper Phillips curves (higher λ_i , hence higher r_i), so a given marginal cost deviation translates into a larger price response. The cost-cascade channel operates through $\widehat{mc}_i$: sector i 's marginal cost depends on its suppliers' prices, which are themselves sticky, so the marginal cost response varies across sectors depending on their IO linkages.

Proposition 4 (Sign of the cost-cascade channel). *Hold flexibility homogeneous at $\bar{r}$ and consider the impact price response to a nominal demand impulse $\Delta \hat{w}$. To first order in $\bar{r}$, the cross-sectional regression slope of the price response on upstreamness attributable to the IO cost-cascade channel is*

$$\frac{\partial}{\partial U} \hat{p}_i^{cost} = \bar{r} \Delta \hat{w} \frac{\text{Cov}(\alpha_i, U_i)}{\text{Var}(U_i)} + O(\bar{r}^2). \quad (14)$$

Hence the cost-cascade channel reinforces the flexibility channel if $\text{Cov}(\alpha_i, U_i) > 0$ and offsets it if $\text{Cov}(\alpha_i, U_i) < 0$.

The result follows from the Neumann expansion $\bar{r}(\mathbf{I} - \bar{r}\Gamma)^{-1}\alpha = \bar{r}\alpha + \bar{r}^2\Gamma\alpha + \dots$. The leading term is $\bar{r}\alpha_i$, whose cross-sectional slope on U_i is $\bar{r} \text{Cov}(\alpha, U) / \text{Var}(U)$. With homogeneous rigidity, a sector's price response is governed by how directly its marginal cost moves with the wage—its labor share α_i —so the cost-cascade channel raises upstream responses relative to downstream ones only if upstream sectors are more labor intensive.

2.4 Special cases and intuition

To build intuition for the general results, consider a two-sector economy ($N = 2$) with a simple chain structure: sector 1 (upstream) sells only to sector 2 (downstream), which sells only to final

consumers. Let γ denote the cost share of sector-1 inputs in sector 2's production, so the IO matrix is

$$\Gamma = \begin{pmatrix} 0 & 0 \\ \gamma & 0 \end{pmatrix}. \quad (15)$$

Sector 1 uses only labor ($\alpha_1 = 1$), while sector 2 uses labor and sector-1 output ($\alpha_2 = 1 - \gamma$). The output-allocation matrix has $a_{12} = \gamma$ (the share of sector 1's output sold to sector 2) and all other entries zero, so the Ghosh inverse is $(\mathbf{I} - \mathbf{A})^{-1} = \begin{pmatrix} 1 & \gamma \\ 0 & 1 \end{pmatrix}$ and the upstreamness values are $U_1 = 1 + \gamma$ and $U_2 = 1$.¹⁴

In this two-sector chain, the system of Phillips curves reduces to:

$$\pi_{1t} = \beta \mathbb{E}_t \pi_{1,t+1} + \lambda_1 (\hat{w}_t - \hat{p}_{1t}), \quad (16)$$

$$\pi_{2t} = \beta \mathbb{E}_t \pi_{2,t+1} + \lambda_2 [(1 - \gamma)\hat{w}_t + \gamma\hat{p}_{1t} - \hat{p}_{2t}]. \quad (17)$$

Equation 16 is a standard single-sector Phillips curve for the upstream industry. Equation 17 shows that downstream inflation depends on both the wage and the upstream price $\hat{p}_{1t}$, the network linkage. Two distinct forces shape the relative price response of the two sectors, and it is useful to separate them at the outset.

The flexibility force is the primary driver. Because sector 1 sells to another firm rather than to consumers, its prices are renegotiated more frequently (Assumption 1), so $\lambda_1 > \lambda_2$. A sector with a steeper Phillips curve passes a larger share of any marginal cost change into its price. This force makes the upstream sector respond more, and it operates regardless of the IO structure. The cost-linkage force is secondary. The downstream sector's marginal cost depends on the sticky upstream price $\hat{p}_{1t}$, so a fraction γ of its costs adjusts only as the upstream price moves. Whether this force makes the upstream or the downstream sector respond more depends on the relative labor intensities of the two sectors.

Solving the two-sector system in closed form under the simplifying assumption that both Phillips curves are static ($\beta = 0$) yields the impact responses.¹⁵ Setting $\Delta\hat{w} \equiv \partial\hat{w}_t/\partial\epsilon_t^m$ for brevity, the static Phillips curve for sector 1 gives $\hat{p}_1 = \lambda_1(\Delta\hat{w} - \hat{p}_1)$, so:

$$\frac{\partial\hat{p}_{1t}}{\partial\epsilon_t^m} = \frac{\lambda_1}{1 + \lambda_1} \Delta\hat{w}, \quad (18)$$

$$\frac{\partial\hat{p}_{2t}}{\partial\epsilon_t^m} = \frac{\lambda_2}{1 + \lambda_2} \left[(1 - \gamma) \Delta\hat{w} + \gamma \cdot \frac{\lambda_1}{1 + \lambda_1} \Delta\hat{w} \right]. \quad (19)$$

¹⁴The row sums of the Ghosh inverse give $U_1 = 1 + \gamma > 1 = U_2$ for any $\gamma > 0$. In a longer chain of K sectors with uniform cost share γ , the most upstream sector has $U_1 = 1 + \gamma + \gamma^2 + \dots + \gamma^{K-1} = (1 - \gamma^K)/(1 - \gamma)$.

¹⁵Setting $\beta = 0$ eliminates forward-looking expectations and yields closed-form expressions at the cost of ignoring the intertemporal dimension. The cross-sectional predictions extend to $\beta > 0$; the main effect of forward-looking pricing is to front-load the response—firms anticipate the future cost and demand path and adjust on impact—so that the peak arrives earlier, as is standard in New Keynesian models. The relative ordering of sectoral responses by upstreamness is unaffected.

The term $\lambda_i / (1 + \lambda_i) = r_i$ is the flexibility force, which is larger for the upstream sector under [Assumption 1](#), directly raising sector 1's response. The bracketed term in the downstream response is the cost-linkage force, with the input price $\hat{p}_{1t}$ entering with weight γ and itself attenuated by upstream stickiness.

The labor-share structure of the chain determines the sign of the cost-linkage force. In this stylized chain, the most upstream sector uses only labor ($\alpha_1 = 1$), so its marginal cost moves one-for-one with the wage while the downstream sector's cost is partly the sticky upstream price; the cost-linkage force then reinforces the flexibility force, making the upstream sector respond more. If instead upstream sectors are more intermediate-input intensive and carry lower labor shares, their marginal cost depends more on other sticky prices and less on the flexible wage, so the cost-linkage force works in the opposite direction and partially offsets the flexibility force. The chain is thus useful for seeing how stickiness compounds across stages, but whether the cost-linkage force reinforces or offsets the flexibility force—and hence which dominates the upstreamness differential—depends on the labor-share structure of the network, a quantitative question I take to the calibration in [Subsection 5.2](#).

For the asymmetry results that follow, define the cascade factor $\phi \equiv \gamma r = \gamma \lambda / (1 + \lambda)$, the per-stage rate at which a cost change passes through a sticky downstream price. In a chain with K stages this factor is applied at each node, and as shown in [Proposition 3\(ii\)](#) the upstream-downstream price differential D_K is increasing in K . The general N -sector model replaces the chain with the full IO matrix Γ , but the two forces are the same, namely the flexibility variation across network positions and cost linkages whose sign depends on the labor-share structure.

This cascade factor reveals why asymmetric rigidity is consequential. If λ differs between upward and downward price adjustment—as the micro evidence strongly suggests—then ϕ differs by sign of the shock. In a single sector, the asymmetry is a level effect; in a K -stage chain, the asymmetry is amplified because longer chains compound the difference in cascade factors across more stages. I formalize this insight in the next subsection.

2.5 Asymmetric price rigidity and network amplification

The results so far characterize how network-determined rigidity heterogeneity shapes the cross-section of price responses under symmetric Calvo pricing. The central theoretical contribution of this paper is to show that *asymmetric* price rigidity—the empirical regularity that prices adjust more readily upward than downward—interacts with the production network to generate a qualitatively new prediction: the upstreamness differential in price responses is substantially larger following expansionary shocks than contractionary ones. This prediction is absent from all existing multi-sector NK models ([Pasten et al., 2020](#); [Rubbo, 2023](#)), which maintain symmetric Calvo pricing, and it cannot arise in representative-firm models regardless of the pricing assumption.

The extension is motivated by a large empirical literature documenting that prices adjust asymmetrically. [Peltzman \(2000\)](#) finds that retail prices respond more quickly and fully to cost increases than to cost decreases across a broad range of products. [Nakamura and Steinsson \(2008\)](#) document

that the fraction of price increases exceeds the fraction of price decreases for most product categories, and that the hazard of price adjustment depends on the sign of the gap between the actual and desired price.¹⁶

Asymmetric Calvo pricing I introduce a parsimonious reduced-form extension of the Calvo framework in which the probability of price resetting depends on the direction of the desired adjustment. Let $\hat{p}_{it}^*$ denote the flexible-price optimum for sector i at time t . A firm in sector i can reset its price with probability

$$1 - \theta_i^s, \quad s \in \{+, -\}, \quad (20)$$

where $s = +$ if the desired reset price exceeds the current price ($\hat{p}_{it}^* > \hat{p}_{it}$) and $s = -$ if $\hat{p}_{it}^* < \hat{p}_{it}$. Asymmetric price rigidity corresponds to $\theta_i^- > \theta_i^+$, meaning firms are less likely to lower their price than to raise it. The baseline Calvo model is the special case $\theta_i^+ = \theta_i^- = \theta_i$. This specification assigns a common sign state s_t to all sectors based on the direction of the aggregate shock, which is a simplification as, in practice, idiosyncratic shocks would cause some firms to wish to raise prices while others wish to cut them. The reduced-form approach captures the dominant direction of desired adjustment following a common monetary shock and is best interpreted as characterizing the average behavior of the cross-section conditional on the aggregate state.

State-dependent Phillips curves Under asymmetric Calvo pricing, the slope of the sectoral Phillips curve becomes state-dependent. Define the state-dependent slope parameters

$$\lambda_i^s = \frac{(1 - \theta_i^s)(1 - \beta\theta_i^s)}{\theta_i^s}, \quad s \in \{+, -\}. \quad (21)$$

Since λ_i^s is decreasing in θ_i^s , asymmetric rigidity ($\theta_i^- > \theta_i^+$) implies $\lambda_i^- < \lambda_i^+$, so the Phillips curve is flatter when firms need to cut prices than when they need to raise them.

The sectoral Phillips curve becomes

$$\pi_{it} = \beta \mathbb{E}_t \pi_{i,t+1} + \lambda_i^{s_t} \left[\alpha_i \hat{w}_t + \sum_{k=1}^N \gamma_{ik} \hat{p}_{kt} - \hat{p}_{it} \right], \quad (22)$$

where $s_t \in \{+, -\}$ depends on the sign of the aggregate demand impulse.

¹⁶Part of the asymmetry in price change frequencies documented by [Nakamura and Steinsson \(2008\)](#) may reflect positive trend inflation rather than a structural asymmetry in adjustment costs. The reduced-form specification adopted here is agnostic about the source of the asymmetry: it captures both genuine downward nominal rigidity (as in menu cost models with asymmetric adjustment bands) and the mechanical effect of positive trend inflation on adjustment hazards. [Appendix A-III](#) develops a menu cost model with asymmetric Ss bands that provides a micro-founded interpretation. See also [Cavallo et al. \(2024\)](#) for a state-dependent pricing model with small and large price changes that generates asymmetric adjustment endogenously.

Network amplification of the asymmetry The key insight is that asymmetric rigidity interacts multiplicatively with network structure. Return to the two-sector chain of [Subsection 2.4](#). After an expansionary shock ($\varepsilon_t^m < 0$), demand rises and firms wish to raise prices. Both sectors face the “upward” slope λ_i^+ , which is relatively steep. The upstream sector’s price increases quickly, and this increase cascades to the downstream sector through input costs. At each stage, the Phillips curve slope is λ^+ , so the cascade operates at full force.

After a contractionary shock ($\varepsilon_t^m > 0$), demand falls and firms wish to cut prices. Now both sectors face the “downward” slope $\lambda_i^- < \lambda_i^+$. The cascade is attenuated at every stage. Even the upstream sector adjusts more slowly, and the downstream sector’s input costs move less, compounding the attenuation through the production chain.

The following result formalizes this logic for an arbitrary production network, before specializing to the chain. Let $r_i^s \equiv \lambda_i^s / (1 + \lambda_i^s)$ denote the sign-dependent price-adjustment probability and $\mathbf{R}^s \equiv \text{diag}(r_1^s, \dots, r_N^s)$, for $s \in \{+, -\}$.

Proposition 5 (Sign-dependent network propagation). *In the static economy ($\beta = 0$), the impact response of the sectoral price vector to a nominal demand impulse of sign s is*

$$\hat{\mathbf{p}}^s = \mathbf{M}^s \boldsymbol{\alpha} \Delta \hat{w}, \quad \mathbf{M}^s \equiv (\mathbf{I} - \mathbf{R}^s \Gamma)^{-1} \mathbf{R}^s, \quad (23)$$

where $\mathbf{M}^s$ is the sign-dependent dampened Leontief matrix. This matrix satisfies:

- (i) **Nesting.** If $r_i^+ = r_i^-$ for all i (symmetric rigidity), then $\mathbf{M}^+ = \mathbf{M}^-$ and price responses are independent of the sign of the shock.
- (ii) **Asymmetry.** If $r_i^+ \geq r_i^-$ for all i , with strict inequality for some i , then $\mathbf{M}^+ \geq \mathbf{M}^-$ elementwise, so $|\hat{p}_i^+| \geq |\hat{p}_i^-|$ for every sector: price responses are uniformly larger following expansionary shocks, and the cross-sectional dispersion of responses with respect to upstreamness is wider.

The proof of part (ii) uses the Neumann series $\mathbf{M}^s = \sum_{k=0}^{\infty} (\mathbf{R}^s \Gamma)^k \mathbf{R}^s$, which converges because $\mathbf{R}^s \Gamma$ is non-negative with spectral radius below one ($r_i^s < 1$ and Γ has row sums below one). Each term is elementwise non-decreasing in every r_i^s , so $\mathbf{R}^+ \geq \mathbf{R}^-$ implies $\mathbf{M}^+ \geq \mathbf{M}^-$ entrywise. The matrix $\mathbf{M}^s$ is the static, sign-dependent analog of the duration-adjusted Leontief sufficient statistic of [Afrouzi and Bhattarai \(2024\)](#). In the symmetric case it collapses to a single propagation matrix governing all responses, as in their framework and in [Rubbo \(2023\)](#) and [Pasten et al. \(2020\)](#); asymmetric rigidity splits it into an expansionary matrix $\mathbf{M}^+$ with shorter effective price durations and a contractionary matrix $\mathbf{M}^-$ with longer ones. The rockets-and-feathers asymmetry is governed by the gap $\mathbf{M}^+ - \mathbf{M}^-$, an object absent from the symmetric multi-sector literature. The chain economy delivers a closed-form characterization of this gap.

Proposition 6 (Asymmetric network amplification: chain economy). *Under asymmetric Calvo pricing with $\theta_i^- > \theta_i^+$ for all i , in a chain of K sectors with uniform input share γ and $\beta = 0$:*

(i) The upstream-downstream price differential under shock sign $s \in \{+, -\}$ is

$$D_K^s = \frac{r^s \gamma (1 - r^s) [1 - (\phi^s)^{K-1}]}{1 - \phi^s}, \quad (24)$$

where $r^s = \lambda^s / (1 + \lambda^s)$, $\phi^s = r^s \gamma$, and the differential is expressed per unit of the wage impulse $\hat{\Delta}w$.

(ii) The ratio of the expansionary to the contractionary differential is

$$\mathcal{R}_K = \frac{r^+(1 - r^+)}{r^-(1 - r^-)} \times \frac{\sum_{j=0}^{K-2} (\phi^+)^j}{\sum_{j=0}^{K-2} (\phi^-)^j}. \quad (25)$$

Under the restriction $r^+, r^- < \frac{1}{2}$ with $r^+ > r^-$: (a) $\mathcal{R}_K > 1$ for all $K \geq 2$; (b) $\mathcal{R}_K$ is strictly increasing in K ; and (c) $\mathcal{R}_K$ converges to the finite limit $\mathcal{R}_\infty = r^+(1 - r^+)(1 - \phi^-) / [r^-(1 - r^-)(1 - \phi^+)]$.

The ratio $\mathcal{R}_K$ has a clean two-factor decomposition. The first factor, $r^+(1 - r^+) / [r^-(1 - r^-)]$, captures the *own-flexibility asymmetry*: the upstream-downstream differential is proportional to $r(1 - r)$, which is larger for r^+ than for r^- when both lie below $\frac{1}{2}$. This factor operates even in a two-sector chain ($K = 2$). The second factor, the ratio of partial geometric sums, captures the *network amplification of the asymmetry*: it equals one at $K = 2$, increases with K , and converges to $(1 - \phi^-) / (1 - \phi^+) > 1$. Longer chains amplify the asymmetry because the stronger expansionary cascade ($\phi^+ > \phi^-$) cumulates over more stages, but the amplification is bounded.

The model thus makes a testable prediction. The interaction between upstreamness and monetary shocks should be driven predominantly by expansionary episodes. Moreover, the asymmetry should be stronger for industries that are more deeply embedded in the production network (higher upstreamness), because the difference in cascade factors cumulates over more stages.¹⁷

3 Data

The empirical analysis combines three data sources: industry-level producer prices, the U.S. IO tables, and a measure of monetary policy stance.¹⁸

3.1 Producer price indices

I use monthly Producer Price Index (PPI) data from the Bureau of Labor Statistics. The baseline sample covers all industries with BLS Producer Price Index coverage at the six-digit NAICS level

¹⁷Beyond the cross-section, the same matrix system has an aggregate implication. The production network flattens the economy-wide Phillips curve, with the slope a decreasing function of network density. Because this paper's contribution is cross-sectional, I develop this result—which complements [Rubbo \(2023\)](#)—in [Appendix A-II](#).

¹⁸Additional details on data construction are in [Appendix A-IV](#).

over the period February 1998 to December 2024. This includes manufacturing (NAICS 311–339), mining, utilities, wholesale trade, retail trade, transportation, information, and selected service sectors. The panel is unbalanced: some PPI series (e.g., Food Manufacturing, NAICS 311) are available from the late 1990s, while others (e.g., Wholesale Trade subsectors) begin around 2003 when the BLS expanded PPI coverage.

Table 1: Industries in the sample and network position

Industry		Upstreamness	Centrality
<i>Q1 (Most downstream) (N = 69)</i>			
Min	Direct life insurance carriers	1.00	0.00
Max	Fan, blower, air purification equipment mfg	1.03	0.00
<i>Q2 (N = 68)</i>			
Min	Cut stone and stone product manufacturing	1.03	0.00
Max	Clay product and refractory manufacturing	1.11	0.00
<i>Q3 (N = 68)</i>			
Min	Dry, condensed, and evaporated dairy products	1.11	0.00
Max	Architectural & structural metals mfg	1.28	0.01
<i>Q4 (Most upstream) (N = 69)</i>			
Min	Portfolio management and investment advice	1.28	0.06
Max	Iron and steel mills and ferroalloy mfg	3.77	0.08
Total industries: 274		1.23	0.02

Notes: For each upstreamness quartile, the table reports the industries with the minimum and maximum upstreamness values. Upstreamness is computed following [Antràs et al. \(2012\)](#) from the BEA Detailed IO tables. Centrality refers to eigenvector centrality, computed as the dominant eigenvector of the direct requirements matrix; it captures an industry’s importance based on the centrality of the industries it supplies. N is the number of industries in the quartile. The full sample includes all industries with BLS PPI coverage.

3.2 IO tables and network position

I construct the production network using the BEA 2012 Detailed IO tables.¹⁹ Specifically, I use the Detailed-level Use table, which records the dollar value of each commodity used as an intermediate input by each industry at the six-digit NAICS level. I match the BEA Detailed industries to the six-digit NAICS industries in the PPI sample using a crosswalk that maps BEA commodity codes to individual six-digit NAICS codes. After matching, I normalize the rows to obtain the technical coefficients matrix $\Gamma = [\gamma_{ik}]$, where γ_{ik} is the cost share of inputs from sector k in sector i ’s total intermediate expenditure.

¹⁹The Detailed table is constructed from the Economic Census and provides the most granular picture of inter-industry flows. I use the 2012 benchmark table throughout the baseline analysis. The BEA also publishes summary-level Use tables for other years; I verify in the robustness section that results are not sensitive to using the 2007 table instead, consistent with the well-documented stability of the U.S. IO structure at the level of inter-sector cost shares ([Carvalho and Tahbaz-Salehi, 2019](#)).

Following [Antràs et al. \(2012\)](#), I define the upstreamness of industry i as

$$U_i = 1 \times \frac{F_i}{Y_i} + 2 \times \frac{\sum_k a_{ik} F_k}{Y_i} + 3 \times \frac{\sum_k \sum_j a_{ij} a_{jk} F_k}{Y_i} + \dots = \mathbf{e}_i' (\mathbf{I} - \mathbf{A})^{-1} \mathbf{1}, \quad (26)$$

where F_i is final demand for industry i 's output, Y_i is total output, and $\mathbf{A}$ is the output-allocation matrix with $a_{ik} = Z_{ik}/Y_i$ (the share of industry i 's output sold to industry k).²⁰ Hence, an industry that sells entirely to final consumers has $U_i = 1$; higher values indicate more intermediate steps to final demand. In the full sample of industries with BLS PPI coverage, upstreamness ranges from 1.00 to 3.77, with substantial variation in production chain position, as shown in [Table 1](#). As an alternative measure of network position, I compute the eigenvector centrality of each industry from the IO matrix. The centrality v_i is proportional to the weighted sum of the centralities of the industries it supplies: $v_i = \frac{1}{\lambda_{\max}} \sum_k \gamma_{ki} v_k$, where $\lambda_{\max}$ is the largest eigenvalue of Γ' . An additional raw measure of network position is the share of intermediate inputs in total purchases, which I use later in robustness tests.

[Figure 1](#) summarizes the production network and the cross-sectional variation in network position. Panel (a) visualizes the network aggregated to the three-digit NAICS level for legibility. Node size is proportional to eigenvector centrality, color indicates the upstreamness quartile (lighter shades for downstream, darker for upstream), and directed edges point from supplier to customer. Two features stand out. First, the network is dense: most sectors are connected to multiple suppliers and customers, generating the inter-industry cost linkages through which monetary shocks propagate. Second, while node centrality and upstreamness are positively correlated, such correlation is imperfect, warranting some robustness analysis.


Panel (b) shows the cross-sectional distribution of upstreamness. The distribution is right-skewed, with a modal value near 1.0 and a long right tail extending beyond 3.5. The right skew implies that a relatively small number of industries sit at the upstream end of the production chain (e.g., basic chemicals, primary metals), while the majority are positioned closer to final demand. The quartile boundaries—marked by dashed vertical lines—divide the sample into four groups with meaningfully different positions in the production chain. The substantial spread between the first and fourth quartiles provides ample cross-sectional variation to identify how network position shapes the transmission of monetary policy shocks.

3.3 Monetary policy shocks

My preferred monetary policy shock measure is the orthogonalized monetary policy surprise (MPS) of [Bauer and Swanson \(2023\)](#), which is available up to December 2023. These shocks are constructed in two steps. First, high-frequency changes in federal funds futures prices in a narrow window around FOMC announcements isolate the market's revision of the expected path of pol-

²⁰The output-allocation matrix $\mathbf{A}$ differs from the cost-share matrix Γ introduced in the model section, which captures the demand side of inter-industry flows. The upstreamness measure requires the supply-side matrix $\mathbf{A}$ because it traces the forward linkages from an industry's output toward final demand. The Leontief inverse $(\mathbf{I} - \mathbf{A})^{-1}$ is sometimes called the output inverse or Ghosh inverse.

Figure 1: The U.S. production network and upstreamness distribution



Notes: Panel (a): Each node represents an industry aggregated to the three-digit NAICS level. Node size is proportional to eigenvector centrality; color indicates upstreamness quartile (lighter = more downstream, darker = more upstream). A directed edge from industry i to industry j indicates that i supplies intermediate inputs to j ; edge width is proportional to the value of the flow from the BEA Detailed IO tables. Panel (b): Cross-sectional distribution of upstreamness across all industries in the sample, computed from BEA Detailed IO tables following [Antràs et al. \(2012\)](#); vertical dashed lines indicate quartile boundaries.

icy rates. Second, the raw surprises are orthogonalized with respect to a set of macroeconomic and financial variables to purge the “Fed information effect”—the concern that surprise rate changes partly reflect the Fed’s superior information about the state of the economy rather than exogenous shifts in the policy rule. The resulting MPS series is available at the FOMC-meeting frequency. I aggregate to monthly frequency by summing surprises within each calendar month (months without FOMC meetings receive a shock of zero).²¹

As a robustness check, I also use the monetary policy shocks of [Jarociński and Karadi \(2020\)](#), which is available up to December 2024. Their approach applies a sign-restriction identification scheme to a high-frequency VAR of interest rate and stock price changes around FOMC announcements. The key identifying assumption is that a pure monetary policy shock moves interest rates and stock prices in opposite directions (higher rates, lower stocks), whereas a central bank information shock moves them in the same direction (higher rates, higher stocks, because the rate increase reflects good economic news). The median-rotation estimator from this decomposition isolates the monetary policy component, purging the information effect through an alternative channel to the Bauer-Swanson orthogonalization.²²

²¹In months with two FOMC meetings (e.g., January and March in some years), the monthly shock is the sum of both surprises. In the sample, approximately 8% of months contain two meetings.

²²Figure A2 of [Appendix A-IV](#) show the time series of the shocks.

3.4 Sample construction and descriptive statistics

The estimation sample is a monthly panel of 271 industries observed from February 1998 through December 2023. It gets extended to end-2024 when using the monetary policy shocks of [Jarociński and Karadi \(2020\)](#). Of the 274 industries in the BEA IO tables, three are excluded for having less than 10 years of PPI data.²³ [Table A2 of Appendix A-IV](#) reports summary statistics. Monthly PPI inflation averages 0.16% with a standard deviation of 1.8%, and industries display substantial heterogeneity in upstreamness.

4 Empirical strategy

This section presents the estimation approach, the identification assumptions and their plausibility, and the interpretation of estimates.

4.1 Estimation

The empirical strategy exploits the interaction of time-series variation in monetary policy shocks with cross-sectional variation in production network position. Specifically, I estimate impulse response functions using the local projection method of [Jordà \(2005\)](#). For each horizon $h = 0, 1, \dots, H$ (with $H = 36$ months), I estimate:

$$\Delta^h p_{it} = \alpha_i^h + \mu_{s(i),t}^h + \beta^h (U_i \times \varepsilon_t^m) + \rho^h \Delta p_{i,t-1} + u_{it}^h, \quad (27)$$

where $\Delta^h p_{it} = \ln(\text{PPI}_{i,t+h}) - \ln(\text{PPI}_{it})$ is the cumulative log price change in six-digit NAICS industry i from month t to $t + h$ (at $h = 0$, the dependent variable is the contemporaneous monthly PPI inflation π_{it}), U_i is the upstreamness of industry i , ε_t^m is the Bauer-Swanson monetary policy shock, α_i^h are the six-digit NAICS industry fixed effects, $\mu_{s(i),t}^h$ are three-digit NAICS sector $\times$ month fixed effects, and $\Delta p_{i,t-1}$ denotes the lagged PPI inflation.^{24,25}

The coefficient of interest is β^h , which captures how the cumulative price response to a monetary shock varies with a sector's position in the production network. Under the model's prediction ([Proposition 3](#)), such coefficient is negative at medium-to-long horizons—meaning that more up-

²³Relaxing the threshold to 60 months adds three additional industries without materially affecting the estimates.

²⁴The main effects of U_i and ε_t^m are absorbed by the industry and sector $\times$ time fixed effects, respectively, and are therefore omitted from the written specification.

²⁵The choice of local projections over vector autoregressions (VARs) is motivated by several considerations. First, [Plagborg-Møller and Wolf \(2021\)](#) establish that local projection and VAR impulse responses estimate the same causal estimand under general conditions—the local projection-VAR equivalence result—and show that the equivalence holds exactly when the local projection lag length equals the VAR lag length. With one PPI lag in the present specification, the local projection estimates are numerically equivalent to those from a one-lag panel VAR with the same interaction structure. Second, local projections are robust to misspecification of the lag structure, which is particularly relevant in a panel setting where the optimal lag length may vary across industries. Third, local projections directly accommodate the interaction term $U_i \times \varepsilon_t^m$ without requiring a nonlinear VAR specification. Fourth, local projections produce estimates at each horizon h independently, so misspecification at one horizon does not contaminate estimates at other horizons.

stream sectors exhibit larger cumulative price declines following a contractionary monetary policy shock, consistent with the amplification mechanism operating through multi-stage production chains. The horizon choice of $H = 36$ months is guided by the model's prediction that the full price adjustment takes several years, while balancing the loss of observations at long horizons in the unbalanced panel. The industry fixed effects absorb any time-invariant differences in average price growth across industries. The sector $\times$ time fixed effects $\mu_{s(i),t}^h$ absorb all time-varying shocks common to industries within the same three-digit NAICS sector. The monetary shock effect ε_t^m is absorbed by the sector $\times$ time fixed effects, so β^h is identified purely from within-sector variation in upstreamness across six-digit industries interacted with the time-series monetary shock variation.²⁶

Identification The identification strategy rests on two assumptions. The first is the exogeneity of monetary policy shocks. That is, conditional on industry fixed effects and lagged macroeconomic controls, the interaction $U_i \times \varepsilon_t^m$ is uncorrelated with the error term. The second is the stability of network position, namely that the upstreamness U_i of each industry is stable over the sample period, reflecting technological characteristics of the industry's product rather than endogenous responses to monetary conditions. These assumptions have several distinct components, each of which I now discuss.

The key requirement for identification is that the FOMC does not target sectoral PPI developments. This is plausible because the FOMC targets aggregate inflation (CPI or PCE), not sector-specific producer prices. Moreover, the Bauer-Swanson shocks are constructed from high-frequency changes in federal funds futures around FOMC announcements, then orthogonalized with respect to macroeconomic and financial variables. By construction, the shock is the component of the policy surprise that is orthogonal to both the systematic policy rule and the market's inference about the Fed's private economic information. Nakamura and Steinsson (2018) demonstrate that the information effect is quantitatively important, and the Bauer-Swanson orthogonalization directly addresses this concern by removing the component of the surprise that is correlated with the Fed's private information about economic conditions.

Even if the monetary policy shocks are exogenous, however, identification of β^h requires that the interaction $U_i \times \varepsilon_t^m$ be exogenous. The main threat would be if the monetary shock were systematically correlated with shocks that differentially affect upstream versus downstream industries. For example, if oil price shocks simultaneously caused the Fed to tighten policy and disproportionately affected upstream commodity-producing sectors, the interaction term would be endogenous. The sector $\times$ time fixed effects directly address this concern by absorbing all time-varying shocks that are common within three-digit NAICS sectors, including for example sectoral demand shifts, commodity price movements, and trade shocks that could potentially correlate

²⁶A variance decomposition indicates that 56% of the total cross-sectional variance in upstreamness is within three-digit sectors, with a median within-sector standard deviation of 0.15 and a median within-sector range of 0.40 (compared to an overall standard deviation of 0.35). The identification is therefore broad-based rather than driven by a handful of within-sector comparisons.

with monetary policy actions. The orthogonalization in the Bauer-Swanson shock provides additional protection, as it removes the correlation between the policy surprise and contemporaneous macroeconomic conditions.

The network stability assumption requires that upstreamness reflects technological characteristics rather than endogenous responses to monetary conditions. This is supported by the fact that the IO structure of an industry—who its customers and suppliers are, and in what proportions—is determined by the physical nature of its product and by long-standing supply relationships. [Antràs et al. \(2012\)](#) document that industry-level upstreamness rankings are highly stable over time. The rank correlation of industry-level upstreamness between the BEA annual IO tables for 2002 and 2022 is 0.94. I use the 2012 BEA Detailed table, which postdates the start of the sample, further reducing concerns about endogeneity.²⁷

Three potential threats deserve explicit discussion. First, if upstream and downstream industries face systematically different demand elasticities, the differential price response could reflect demand rather than network structure. I address this by controlling for industry fixed effects and focusing on the interaction term between upstreamness and monetary policy shocks. Second, the estimates could be biased if industry demand, productivity, or cost shocks are correlated with both the monetary shock and upstreamness. The Bauer-Swanson orthogonalization mitigates this concern by removing the endogenous component of the policy surprise. Also, the sector-time fixed effects absorb time-varying factors that are common to the industries in the same sector. Yet, I recognize that time-varying industry-specific demand shocks remain a potential confound. Third, while the manufacturing cross-section ($N = 209$) is substantially larger than samples common in the literature, concerns about influential observations remain; I address this in the robustness section by examining outlier sensitivity.

Relationship with the model The coefficient β^h in [Equation 27](#) measures how the cumulative price response to a monetary shock varies across industries with different upstreamness values. [Proposition 3](#) predicts a negative coefficient at medium-to-long horizons, with two forces shaping its magnitude. The first, and quantitatively dominant, is the systematic relationship between network position and intrinsic price rigidity. Upstream sectors sell to other firms, and B2B prices adjust more frequently than consumer-facing prices. This generates a correlation between upstreamness and the Calvo parameter, such that more upstream sectors are inherently more flexible. The second force operates through the cost-cascade channel, whereby each sector’s marginal cost depends on its suppliers’ prices, which are themselves sticky, introducing an additional source of

²⁷ [Atalay \(2017\)](#) document that while the identity of an industry’s top suppliers is remarkably stable over decades, the intensity of linkages can shift. The high rank correlation in upstreamness, however, suggests that the stable supplier identity is the dominant force for the network position measure used here. [Ghassibe \(2025\)](#) develops a model in which the number of supplier relationships is itself endogenous and procyclical, so that network density varies over the business cycle. While this mechanism could in principle generate time variation in the strength of network amplification, it does not threaten the cross-sectional identification: the sector $\times$ time fixed effects absorb any aggregate or sector-level variation in network density, and the estimates exploit only within-sector-month differences in upstreamness rankings, which are stable over time.

cross-sectional variation.

Two features of the theory-to-empirics mapping deserve comment. First, the estimating [Equation 27](#) is a reduced-form specification, not a structural equation derived from the model. The model motivates the choice of regressors—the interaction between upstreamness and the monetary shock—and delivers qualitative predictions about the sign and horizon profile of β^h , but the local projection does not impose the model’s cross-equation restrictions. This reduced-form approach has the advantage of being robust to misspecification of the model’s specific functional forms (Calvo vs. menu cost, log-linear vs. nonlinear) while still providing a direct test of the model’s cross-sectional prediction. Second, the model is cast in terms of gross-output prices, while the empirical literature on monetary transmission typically studies CPI or GDP deflator responses. The PPI data used here are a natural empirical counterpart to the model’s gross-output price indices, as they measure the price of a sector’s output inclusive of intermediate input costs, which is precisely the object that the sectoral Phillips curve in [Equation 6](#) describes. A value-added price index—which strips out intermediate input costs—would be the appropriate measure for studying sectoral real value added, but would lose the cost-cascade channel that the model emphasizes.²⁸

4.2 Asymmetric effects of expansionary and contractionary shocks

The baseline specification constrains the interaction coefficient β^h to be the same for positive and negative monetary policy shocks. However, the model’s mechanism—network-determined flexibility interacting with asymmetric price adjustment—may generate different cross-sectional patterns depending on the direction of the shock. Asymmetric price rigidity implies that price increases propagate more readily than price decreases: when the Fed cuts rates and demand rises, upstream firms—being more flexible—raise prices readily, and the upstreamness differential is large. When the Fed tightens and demand contracts, the flatter downward Phillips curve slope attenuates the cross-sectional differential.

To test for asymmetry, I augment [Equation 27](#) with indicator variables for the sign of the shock. Define $D_t^+ = \mathbf{1}\{\varepsilon_t^m > 0\}$ to identify months in which monetary policy was predominantly contractionary and $D_t^- = \mathbf{1}\{\varepsilon_t^m < 0\}$ to identify months in which it was mainly expansionary. Then, for each horizon $h = 0, 1, \dots, H$, I estimate the local projection:

$$\Delta^h p_{it} = \alpha_i^h + \mu_{s(i),t}^h + \beta^{h,+} (U_i \times \varepsilon_t^m \times D_t^+) + \beta^{h,-} (U_i \times \varepsilon_t^m \times D_t^-) + \rho^h \Delta p_{i,t-1} + u_{it}^h, \quad (28)$$

The key coefficients are $\beta^{h,+}$, which captures the upstreamness effect on price responses during contractionary episodes, and $\beta^{h,-}$, which captures the effect during expansionary episodes.²⁹ Un-

²⁸A subtlety in the theory-to-empirics mapping concerns the functional form. The local projection specification in [Equation 27](#) uses U_i linearly, consistent with the theoretical prediction. As a robustness check, I replace U_i with $\log(U_i)$.

²⁹Over the baseline sample, 118 months feature contractionary shocks ($\varepsilon_t^m > 0$) and 93 months feature expansionary shocks ($\varepsilon_t^m < 0$), with comparable mean magnitudes (0.040 and -0.044 , respectively). An additional 100 months have zero shock (no FOMC meeting). As in [Equation 27](#), all lower-order terms ($U_i, \varepsilon_t^m \times D_t^\pm, D_t^\pm$) are absorbed by the

der the null of symmetric effects, $\beta^{h,+} = \beta^{h,-}$ at all horizons.

5 Results

This section presents the empirical findings. I first report the baseline interaction estimates. Second, I discuss the link between the empirical estimates and the model. Third, I examine asymmetric effects across expansionary and contractionary shocks. Fourth, I provide evidence of the real activity implications. Finally, I present auxiliary evidence that costs propagate through buyer-seller linkages in the data.

5.1 Network position and price dynamics

I first present the results about the differential effect of a monetary policy shock according to the network position of a given industry. Panel (a) of [Figure 2](#) shows the estimated impulse response of the interaction coefficient $\hat{\beta}_t$ of [Equation 27](#), tracing out how the differential price response between upstream and downstream industries evolves over time following a monetary policy shock.³⁰ I find that following a contractionary monetary shock, industries with higher upstreamness experience larger cumulative price declines. The interaction coefficient is near zero at impact, becomes progressively more negative and significant, reaching the peak one year and a half after the shock. The effect begins to dissipate after 27 months, consistent with eventual full pass-through across all sectors.³¹

These results are consistent with upstream prices being more flexible—or equivalently, downstream prices being stickier due to the accumulation of nominal rigidities across successive production stages. The magnitude buildup is in line with the model’s prediction that the rigidity differential across network positions cumulates over time as more flexible upstream sectors progressively complete their price adjustment while downstream sectors remain sticky. The hump-shaped pattern matches the theoretical prediction from [Proposition 3](#): the upstreamness effect is transient, operating precisely in the medium-run window where nominal rigidities bind differentially across network positions, before long-run price flexibility restores uniform adjustment.³²

Panel (b) presents impulse response functions estimated separately for each upstreamness quartile. The four quartiles display a monotonic ordering in their price responses. Industries in the most upstream quartile (Q4)—which includes petroleum refining, basic chemicals, iron and steel mills, and other primary materials producers—show the largest cumulative price adjustment, while the most downstream quartiles (Q1 and Q2)—which include food manufacturing,

industry and sector $\times$ time fixed effects and are omitted from the written specification.


³⁰[Table A3](#) of [Appendix A-V](#) reports the full set of local projection estimates.

³¹The loss of significance at long horizons reflects both the dissipation of the real effect (as prices converge to their flexible-price levels) and the mechanical increase in estimation uncertainty: longer-horizon local projections use shorter effective samples in the unbalanced panel.

³²The hump shape is also consistent with models of information rigidity ([Mankiw and Reis, 2002](#)), where firms update their information sets infrequently. In such models, the full effect of a shock on prices builds gradually as more firms observe and respond to the changed monetary conditions.

apparel, furniture, and retail-facing sectors—show responses that are small and statistically indistinguishable from zero at all horizons. The pattern— $Q4 > Q3 > Q2 \approx Q1 \approx 0$ —is precisely what the model predicts.³³ The differential timing across quartiles is informative. The most upstream industries begin showing cumulative responses around the 6-month horizon, while downstream industries remain flat. The gap between Q4 and Q1 widens progressively through 12 months before narrowing at longer horizons, consistent with eventual full pass-through.

Figure 2: Upstreamness and monetary policy transmission



Notes: Panel (a) plots the coefficient on the interaction between upstreamness and monetary policy shocks from Equation 27. Panel (b) plots the cumulative price response to a one-pp monetary shock for industries in each upstreamness quartile, estimated from separate local projections. Shaded bands are 90% confidence intervals based on Driscoll-Kraay standard errors.

The interpretation of the coefficients in economic magnitudes is as follows. The peak interaction coefficient is -0.113 one year and a half after the shock, and the interquartile range of upstreamness is approximately 0.25 units. Thus, per unit of the monetary shock, an industry at the 75th percentile of upstreamness exhibits a peak cumulative price response that is about 2.8 percentage points larger than an industry at the 25th percentile. The quartile estimates provide an even sharper comparison: the most upstream quartile (Q4) shows large and significant cumulative responses, while the most downstream quartile (Q1) is statistically indistinguishable from zero.

5.2 Quantitative fit

Having established the baseline empirical facts, I now ask whether the model from Section 2 can qualitatively capture the cross-sectional pattern. I calibrate the multi-sector model using the actual

³³The near-zero response of Q1 and Q2 industries does not imply that their prices are perfectly rigid; rather, their cumulative response to the interaction with upstreamness is small because they sit close to final demand and face fewer intermediate pricing stages.

BEA IO matrix and simulate the implied path of interaction coefficients $\beta^{h,\text{model}}$ for comparison with the estimated $\hat{\beta}^h$. The calibration has two pricing parameters and a demand path. First, $\bar{\theta}$ is the average Calvo parameter governing the baseline frequency of price adjustment. Second, δ captures the cross-sectional variation in price rigidity with respect to network position $\theta_i = \bar{\theta} + \delta(U_i - \bar{U})$, where $\bar{U}$ is the mean upstreamness. A negative δ means more upstream sectors have more flexible prices, consistent with [Assumption 1](#). Third, the model must be fed a path for the nominal demand impulse $\hat{w}_t$, which is the channel through which the monetary shock reaches sectoral marginal costs.

I discipline both $\bar{\theta}$ and δ externally, using the sectoral price-adjustment frequencies of [Pasten et al. \(2020\)](#), computed from the confidential micro data underlying the BLS Producer Price Index for 341 sectors. Mapping these frequencies to the upstreamness measure, I regress each sector's monthly frequency of price adjustment on its upstreamness. The slope is positive and statistically significant, implying $\hat{\delta} = -0.152$, which means that more upstream sectors reprice more frequently. The average monthly frequency of 14.8% implies $\hat{\theta} = 0.85$, a mean price duration of about seven months, in the range of existing producer-price estimates ([Nakamura and Steinsson, 2008](#)). Reassuringly, the same regression run on my own PPI panel yields a nearly identical slope ($\hat{\delta} = -0.166$), confirming that the flexibility–upstreamness relationship is robust across the standard external dataset and the panel used in the empirical analysis.³⁴

For the demand impulse, rather than impose a monotone decay, I discipline its path to the empirical transmission of the monetary shock to real activity. A monetary shock affects demand with a hump-shaped, delayed profile. Estimating the cumulative response of aggregate industrial production to the Bauer-Swanson shock in my sample yields a trough at seventeen months, consistent with the well-documented hump-shaped output response to monetary shocks (e.g., [Christiano et al., 2005](#); [Romer and Romer, 2004](#)). I therefore feed the model a demand impulse $\hat{w}_t$ with this empirically-disciplined shape (a smooth hump troughing at seventeen months). This matters because the price response inherits the timing of the demand impulse that drives it. That is, a monotone-decaying impulse mechanically produces an early price peak, whereas the empirically hump-shaped impulse produces a delayed one. The demand path is a feature of the monetary-transmission block (the Euler equation and Taylor rule with interest-rate smoothing in [Subsection 2.1](#)), not of the cross-sectional pricing mechanism.

[Figure 3](#) presents the fit and the decomposition of the upstreamness effect. Panel (a) plots the model-implied interaction coefficient against the empirical estimates. The model reproduces all three salient features of the data: the near-zero response on impact, the gradual build-up to a peak at eighteen months, and the subsequent reversion. The model peak of -0.105 captures 93% of the empirical peak of -0.113 , with both peaking at roughly one and a half years. The near-zero

³⁴The relationship is systematic but not deterministic. As [Pasten et al. \(2020\)](#) document, some upstream sectors are among the stickiest (semiconductor manufacturing, with a duration above four years) while some downstream sectors are among the most flexible (food). The slope δ captures the central tendency—upstream sectors are more flexible *on average* because they sell predominantly to other firms—not a one-to-one mapping. A measure based on large ($> 2\%$) price changes in my own panel is more contaminated by volatility (its correlation with PPI volatility is 0.93), which is why I rely on the [Pasten et al. \(2020\)](#) frequencies for the baseline calibration.

impact response is itself informative. It arises because the demand impulse is small on impact and builds over time, so that early on there is little for either upstream or downstream sectors to respond to. A model driven by an impulse that peaks on impact would instead counterfactually generate its largest cross-sectional dispersion at short horizons. That an externally-disciplined model—with the rigidity differential measured from independent sectoral frequencies and the demand path measured from the output response—reproduces the sign, timing, and magnitude of the cross-sectional response without fitting any parameter to it provides a credible link between the theoretical mechanism and the reduced-form evidence.

Panel (b) turns to the mechanism, decomposing the model-implied interaction coefficient into the two channels through which the network can operate. The flexibility channel—the systematic variation in price rigidity across network positions—is isolated by keeping the heterogeneous θ_i but removing IO linkages ($\Gamma = \mathbf{0}$); the cost-cascade channel—the propagation of cost shocks along the production structure—is isolated by retaining Γ but homogenizing rigidity ($\theta_i = \bar{\theta}$). At the peak horizon the flexibility channel accounts for roughly 120% of the total interaction coefficient, while the cost-cascade channel contributes about -50% . This means that holding rigidity constant, the IO network attenuates the upstreamness differential rather than amplifying it, with the interaction of the two channels making up the remaining $+30\%$. This sign is not an artifact of the calibration. By [Proposition 4](#), the cost-cascade channel’s cross-sectional contribution is proportional to $\text{Cov}(\alpha_i, U_i)$, and in U.S. data upstream sectors are intermediate-input intensive rather than labor intensive ($\text{Corr}(U_i, \alpha_i) = -0.90$), so cost linkages work against the flexibility advantage of upstream sectors. The heterogeneous rigidity channel must therefore dominate for the observed pattern to emerge.

This decomposition clarifies the mechanism. The production network determines the cross-sectional distribution of nominal rigidity—network position governs whether a sector sells primarily to other firms or to final consumers, and B2B prices are empirically more flexible. This network-determined rigidity heterogeneity, rather than cost-cascade propagation, is the primary driver of the cross-sectional price response pattern. The network, however, remains essential—without it, there is no structural reason for θ_i to vary systematically with upstreamness—but its role is to shape the distribution of rigidity rather than to propagate price adjustments across stages. This distinction matters for policy, because structural changes that alter the composition of B2B versus B2C activity (e.g., vertical integration, platform-mediated disintermediation) change the distribution of rigidity across the network, with first-order consequences for how monetary policy transmits to prices.


The calibrated model also speaks to the aggregate implication that motivates the paper. Two economies with identical average price rigidity but different production structures display different degrees of monetary non-neutrality. I compare the aggregate (Domar-weighted) price response to a contractionary shock under the network calibration—the full BEA matrix Γ and heterogeneous rigidity $\theta_i = \hat{\theta} + \hat{\delta}(U_i - \bar{U})$ —against a representative-firm benchmark that holds average rigidity fixed but shuts down the network ($\Gamma = \mathbf{0}$, homogeneous $\theta_i = \hat{\theta}$). The network economy

generates modestly more non-neutrality. At the point of widest divergence, around one year after the shock, its cumulative price response is roughly 8% smaller in magnitude, though the two adjust at a similar pace. The gap is modest precisely because, as the decomposition above shows, the IO cost-cascade channel offsets much of the amplification coming from rigidity heterogeneity.

The Calvo framework is analytically convenient but imposes a constant, exogenous probability of price adjustment, missing the nonlinear selection channel captured by menu cost models. That is, a monetary shock pushes some firms past their adjustment thresholds and the resulting price changes can push downstream buyers past their own thresholds, triggering further adjustments. This extensive-margin selection effect is distinct from the intensive-margin cost-cascade channel analyzed in the decomposition above. The cost-cascade channel is partially offsetting in U.S. data because upstream sectors have low labor shares; the extensive-margin selection effect, by contrast, reinforces the flexibility channel because it operates more strongly for sectors that adjust more frequently, which are the upstream ones. [Appendix A-III](#) shows that a multi-sector menu cost model with Ss pricing, strategic complementarities, and IO linkages generates peak interaction coefficient of -0.125 , compared to -0.105 for the externally-calibrated Calvo model and -0.113 in the data. The two pricing frictions bracket the data from below and above, indicating that the cross-sectional pattern is robust to the choice of pricing friction.³⁵ Both time-dependent and state-dependent models generate the hump-shaped interaction between upstreamness and monetary non-neutrality, confirming that the result is driven by the network-determined distribution of rigidity rather than the specific form of nominal rigidity. The menu cost model’s primary role in the paper is not to improve the baseline fit—the Calvo model already matches the data well—but to give the asymmetry of [Subsection 2.5](#) a structural foundation: the sign-dependent rigidity assumed there in reduced form arises endogenously from asymmetric Ss adjustment bands, under which firms are more willing to raise prices than to cut them.

³⁵The menu cost model adopts the generalized Ss framework of [Caballero and Engel \(1999\)](#). Firms adjust when the price gap $|p_{it} - p_{it}^*|$ exceeds a threshold, with potentially asymmetric upper and lower bands. The pricing rule is parameterized rather than derived from a fixed-point problem with lump-sum adjustment costs, which simplifies the multi-sector analysis at the cost of not micro-founding the threshold values. The key parameter—the ratio of upper to lower adjustment bands—is calibrated from micro data on the relative frequency of price increases and decreases.

Figure 3: Quantitative fit and decomposition of the upstreamness effect



Notes: Panel (a) plots the model-implied interaction coefficient between upstreamness and monetary policy shocks from the multi-sector Calvo model with externally calibrated rigidity differential, alongside the empirical estimate with 90% confidence band from Equation 27 ($\bar{\theta} = 0.85$ and $\delta = -0.152$, both disciplined externally from the Pasten *et al.* (2020) sectoral price-adjustment frequencies; the nominal demand impulse is disciplined to the estimated aggregate output response to the monetary shock). Panel (b) decomposes the model-implied interaction coefficient $\hat{\beta}^h$ into “flexibility only”, which keeps the heterogeneous θ_i but sets $\Gamma = 0$; and “cascade only”, which uses homogeneous $\theta_i = \bar{\theta}$ but retains the full IO matrix.

5.3 Asymmetric network amplification


I now turn to examine the differential effects of expansionary and contractionary monetary policy shocks. Figure 4 presents the results of the asymmetric specification in Equation 28. The point estimates suggest an asymmetry consistent with the theoretical prediction of Proposition 6: the upstreamness differential is driven predominantly by expansionary shocks. Panel (a) shows that the expansionary interaction coefficient is negative and significant at medium-to-long horizons, reaching a peak of -0.25 about 27 months after the shock. The effect builds gradually, displaying the same hump-shaped profile as the pooled baseline but with substantially larger point estimates. By contrast, the contractionary interaction coefficient is generally not significant and considerably noisier.³⁶

The pattern at medium horizons is consistent with the model’s prediction. When the Fed eases and demand rises, upstream prices—being more flexible—increase readily. At each stage, the Phillips curve slope is the “upward” parameter λ^+ , which is relatively steep, so the cross-sectional

³⁶Table A3 of Appendix A-V reports the local projection estimates. Wald tests reject symmetry at longer horizons (i.e., after 20 months), but not at earlier ones, where the standard errors on the contractionary coefficient—over 50% larger than on the expansionary side—limit power at individual horizons. To sharpen inference, I average each sign’s interaction coefficient across horizons and test their equality in a single projection of the window-averaged price response. Over both $h = 12\text{--}24$ and $h = 0\text{--}36$, the mean expansionary coefficient is statistically different from the mean contractionary one. The asymmetry is therefore statistically significant once the effects accumulate.

differential between upstream and downstream industries is large. When the Fed tightens and demand falls, each stage faces the flatter “downward” slope $\lambda^- < \lambda^+$, attenuating the differential response. [Proposition 6](#) shows that the ratio of the expansionary to contractionary differential exceeds one and is increasing in the number of production stages, so the asymmetry is amplified along longer supply chains.

Figure 4: Asymmetric effects from expansionary and contractionary monetary shocks



Notes: Panel (a) plots the coefficient on the interaction term between upstreamness and monetary policy shocks with a dummy for whether the shock is expansionary or contractionary from the asymmetric specification in [Equation 28](#). Panel (b) plots each industry’s cumulative price response two years after the shock against its upstreamness, estimated separately on months with expansionary shocks and contractionary shocks; the vertical axis is truncated at ± 1.5 ; three observations fall outside this range (two expansionary, one contractionary). Shaded bands are 90% confidence intervals based on Driscoll-Kraay standard errors.

Panel (b) of [Figure 4](#) provides cross-sectional corroboration. For each industry, I estimate the cumulative 2-year price response to a one-unit monetary shock separately for expansionary and contractionary shocks. During expansionary episodes, the negative relationship between upstreamness and the price response is clearly visible, with more upstream industries exhibiting larger cumulative price changes. During contractionary episodes, the cross-sectional relationship is flatter and noisier, consistent with the panel results in panel (a).

The finding of asymmetry provides micro-founded, cross-sectional evidence for a pattern that has long been documented at the aggregate level: prices rise more readily than they fall. My results show that when price flexibility varies systematically with network position, the cross-sectional dispersion in price responses is larger following expansionary shocks than contractionary ones. The mechanism speaks to the broader macroeconomic puzzle of why disinflation is costly while inflation can accelerate rapidly. Rate cuts generate large cross-sectional price dispersion as more flexible upstream prices increase readily; rate hikes encounter the flatter downward Phillips curve slope, muting the cross-sectional differential. This asymmetry in the transmission mechanism implies that central banks face a harder task when tightening than when easing—a feature

that standard models abstracting from production structure cannot capture.³⁷

The asymmetric estimates can be used to back out implied structural parameters under the model’s restrictions. Under the asymmetric Calvo model, the asymmetry ratio $\mathcal{R}_K$ from [Equation 25](#) decomposes into an own-flexibility factor $r^+(1 - r^+)/[r^-(1 - r^-)]$ and a network factor that grows with the effective chain length $\bar{K}$. The ratio of the expansionary peak coefficient (at $h = 27$) to the pooled peak (at $h = 18$), $|\hat{\beta}_{27}^{h,-}|/|\hat{\beta}_{18}^h| \approx 2.2$, combined with the median upstreamness $\bar{K} \approx 2$, implies an own-flexibility factor of roughly 2.2 (since the network factor equals one at $\bar{K} = 2$). This back-of-the-envelope calculation is illustrative rather than an estimate, since the two coefficients are measured at different horizons. Solving $r^+(1 - r^+)/[r^-(1 - r^-)] = 2.2$ for the ratio of adjustment probabilities suggests that the probability of upward price adjustment is roughly twice the probability of downward adjustment, in line with the central estimates in [Peltzman \(2000\)](#) and [Nakamura and Steinsson \(2008\)](#), who document ratios close to two.

5.4 Real activity implications

The preceding analysis documents that upstream industries exhibit larger cumulative price adjustments following monetary shocks. A natural question is whether this differential price flexibility also shapes the cross-sectional distribution of real activity responses. The New Keynesian framework predicts that industries whose prices adjust more flexibly should absorb monetary shocks primarily through price changes, preserving real output; industries with sticky prices instead absorb shocks through quantity rationing. Because upstream industries have more flexible prices, they should experience smaller output declines following a contractionary monetary shock and smaller output expansions following an easing—a mirror image of the price results.

To test this prediction, I estimate the same local projection interaction specification using cumulative log growth in sector-level industrial production (IP) as the dependent variable:

$$\Delta^h y_{it}^{IP} = \alpha_i^h + \beta^{h,IP} (U_i \times \varepsilon_t^m) + \gamma^h U_i + \delta^h \varepsilon_t^m + \rho^h \Delta y_{i,t-1}^{IP} + u_{it}^h, \quad (29)$$


where $\Delta^h y_{it}^{IP} = \log IP_{i,t+h} - \log IP_{i,t}$ is the cumulative log IP growth from t to $t + h$. Sector-level IP indices are available from FRED for all 20 three-digit NAICS manufacturing sectors (311–339), providing a direct real-side counterpart to the PPI-based price regressions.³⁸ Because IP is measured at the three-digit NAICS level—a coarser classification than the six-digit level used in the price analysis—the sector $\times$ time fixed effects used in the price regressions cannot be included here (they would be collinear with the industry fixed effects). The specification therefore uses industry

³⁷The results may also be relevant for understanding episodes of supply-chain-driven inflation, such as the post-COVID price surge. To the extent that supply disruptions generate cost increases at upstream stages, asymmetric price rigidity implies that these increases propagate through production chains with relatively little attenuation, because upward price adjustment faces less rigidity, while a symmetric unwinding when supply normalizes is attenuated on the downward side. This is consistent with the observed persistence of inflation even after supply bottlenecks eased, though I do not test this channel directly.

³⁸The IP indices are winsorized at the 1st and 99th percentiles within each industry to limit the influence of extreme observations.

fixed effects only, with the monetary shock main effect included as a regressor. The coarser IP classification reduces the cross-sectional sample from 271 to 20 industries and limits the within-sector variation in upstreamness available for identification, so the IP results should be interpreted as providing directional evidence rather than precise estimates. The model predicts that, conditional on the aggregate shock, more upstream industries should exhibit relatively larger (i.e., less negative) output responses, because their greater price flexibility shields real activity from quantity adjustment. I also estimate the asymmetric version of Equation 29, replacing the single interaction with separate expansionary and contractionary terms as in Equation 28.

Figure 5: Effects on industrial production



Notes: Panel (a) plots the coefficient on the interaction coefficient between upstreamness and monetary policy shocks from Equation 29. Panel (b) plots the coefficient on the interaction term between upstreamness and monetary policy shocks with a dummy for whether the shock is expansionary or contractionary. Shaded bands are 90% confidence intervals based on Driscoll-Kraay standard errors. IP indices are winsorized at the 1st and 99th percentiles within each industry.

Figure 5 presents the results. Panel (a) shows the symmetric interaction coefficient for industrial production. The point estimates are positive and statistically significant at medium-to-long horizons. At the peak, the coefficient implies that a manufacturing sector at the 75th percentile of upstreamness experiences a cumulative output decline roughly 1.8 percentage points smaller than one at the 25th percentile, in response to a one percentage point contractionary monetary shock. The direction is consistent with the model's prediction that upstream industries—whose prices adjust more quickly—experience relatively smaller output declines following monetary shocks. The estimates are noisier than the price regressions, which is expected as the IP cross-section contains only 20 three-digit NAICS sectors versus 271 six-digit industries in the price analysis, reducing statistical power.³⁹

Panel (b) decomposes the IP interaction into expansionary and contractionary episodes. The

³⁹Table A4 of Appendix A-V reports the coefficient estimates at all horizons.

contractionary interaction is large and significant one year after the shock. By contrast, the expansionary interaction is uniformly small and insignificant. This pattern is the real-side complement of the price asymmetry in [Subsection 5.3](#), not its repetition. In line with the New Keynesian framework, during contractions upstream industries cut prices more aggressively and, because they adjust through prices, they preserve real output. Downstream industries, whose prices are rigid, instead absorb the shock through quantity rationing—their output falls more. Price flexibility substitutes for quantity adjustment, so cross-sectional dispersion shows up on the price side under expansionary shocks (when prices move readily) and on the output side under contractionary shocks (when downward price rigidity forces output to bear the adjustment), consistent with the asymmetric price rigidity mechanism in the model.

5.5 Supply chain price cascades

I have established that upstream industries exhibit larger cumulative price responses to monetary shocks and that the cross-sectional pattern is primarily driven by network-determined rigidity heterogeneity rather than cost-cascade propagation through IO linkages. Nevertheless, the model predicts that upstream price adjustments do propagate to downstream industries through intermediate input costs, even if this channel is not the primary driver of the upstreamness differential. I provide auxiliary evidence on this cost pass-through channel using a cascade regression that controls for realized supplier prices.⁴⁰

For each industry i and month t , I construct the input-cost-weighted cumulative change in supplier prices:

$$\Delta^h \bar{p}_{it}^{\text{supplier}} = \sum_{k=1}^N \gamma_{ik} \Delta^h p_{kt}, \quad (30)$$

where γ_{ik} is the cost share of inputs from sector k in industry i 's production (from the IO table) and $\Delta^h p_{kt}$ is the cumulative log price change in sector k over horizon h . For each horizon $h = 0, 1, \dots, H$, I estimate the local projection:

$$\Delta^h p_{it} = \alpha_i^h + \mu_{s(i),t}^h + \phi^h \Delta^h \bar{p}_{it}^{\text{supplier}} + \beta^{h,\text{net}} (U_i \times \varepsilon_t^m) + \rho^h \Delta p_{i,t-1} + u_{it}^h. \quad (31)$$


The coefficient ϕ^h captures the pass-through from upstream supplier prices to downstream industry prices—the supply chain cascade. The coefficient $\beta^{h,\text{net}}$ captures any residual effect of network position beyond what operates through the supplier price channel.

[Figure 6](#) visualizes the cascade results.⁴¹ Panel (a) shows a large, positive supplier price coefficient, indicating that upstream price changes pass through to downstream industries via the cost-cascade channel. A one-unit increase in the input-cost-weighted supplier price index raises the

⁴⁰Because the supplier price is itself an outcome of the monetary shock, this regression conditions on a post-treatment mediator rather than a pre-determined control. The attenuation of the upstreamness interaction is therefore suggestive of cost pass-through operating through buyer-seller linkages, not a strict causal decomposition.

⁴¹[Table A5](#) of [Appendix A-V](#) reports the regression results.

Figure 6: Supply chain cascade



Notes: Panel (a) plots the coefficient on supplier prices and the coefficient on the interaction term between upstreamness and monetary policy shocks capturing any residual effect of the network position from the cascade regression in Equation 31. Panel (b) plots the coefficient on supplier prices and the coefficient on the interaction term between upstreamness and monetary policy with a dummy for whether the shock is expansionary or contractionary. Shaded bands are 90% confidence intervals based on Driscoll-Kraay standard errors.

downstream industry's own price by 0.18–0.40 units over the first six months and reaches about 0.74 by eighteen months—substantial but incomplete pass-through, dampened initially by nominal rigidities and reverting somewhat thereafter as the monetary impulse dissipates.⁴² The residual upstreamness interaction—any effect beyond the supplier price channel—remains negative and statistically significant at intermediate horizons, even on the much smaller IO-coverage sample. That a pure input-output construct like the input-cost-weighted supplier price enters strongly and predicts downstream price changes points to genuine network propagation rather than an omitted industry characteristic. At the same time, the upstreamness interaction is not fully absorbed by the supplier price, indicating that the differential does not operate solely through measured input-cost pass-through—consistent with the flexibility channel being its primary driver (Subsection 5.2). The supplier price variable itself reflects both the flexibility channel (more flexible upstream sectors adjust faster) and the cost-cascade channel that transmits those adjustments downstream, so it cannot by itself distinguish between them.

Panel (b) examines the residual interaction in response to expansionary and contractionary monetary policy shocks. Consistent with the asymmetric results from Subsection 5.3, the cascade mechanism is driven almost entirely by expansionary shocks. The expansionary interaction coefficient is large, negative, and significant at horizons beyond 10 months, while the contractionary interaction is statistically insignificant. Together, the two panels of Figure 6 are consistent with costs propagating through buyer-seller linkages, with the propagation concentrated in expansion-

⁴²This gradual convergence is consistent with Nakamura and Steinsson (2008), who document producer-price adjustment every 8–11 months.

ary episodes where upward price flexibility allows rapid cost pass-through.

5.6 Input intensity and network neighbors

The cost-propagation channel predicts two further testable implications. First, the upstreamness effect should be amplified for industries with higher intermediate input shares, since these industries are more exposed to upstream cost pass-through. Second, an industry's price response should depend not only on its own upstreamness but also on the upstreamness of its suppliers—a direct test of network propagation.

To test for input share amplification, I augment the baseline specification with a triple interaction between upstreamness, intermediate input share, and the monetary shock:

$$\Delta^h p_{it} = \alpha_i^h + \mu_{s(i),t}^h + \beta_1 (U_i \times \varepsilon_t^m) + \beta_2 (U_i \times \alpha_i^{II} \times \varepsilon_t^m) + \beta_3 (\alpha_i^{II} \times \varepsilon_t^m) + \rho^h \Delta p_{i,t-1} + u_{it}^h. \quad (32)$$

The specification includes the lower-order interaction $\alpha_i^{II} \times \varepsilon_t^m$ to ensure that $\hat{\beta}_2$ captures the marginal effect of input intensity conditional on the direct effect of intermediate input share on monetary transmission. If cost pass-through operates through buyer-seller linkages, $\hat{\beta}_2$ should be negative, as higher input intensity amplifies the transmission of upstream price changes. Panel (a) of [Figure 7](#) plots the triple interaction coefficient $\hat{\beta}_2$. The estimates are generally negative and statistically significant during the second year after the shock. Industries that are both upstream and input-intensive respond more strongly, and the amplification builds over medium horizons as cost pass-through accumulates through the production chain.


I now turn to the effect of the upstreamness of the suppliers. To test for it, I replace the industry's own upstreamness with the input-cost-weighted average upstreamness of its suppliers:

$$\bar{U}_i^{\text{supplier}} = \sum_k \gamma_{ik} U_k, \quad (33)$$

where γ_{ik} is the share of industry i 's intermediate inputs purchased from industry k . I then estimate the local projection with $\bar{U}_i^{\text{supplier}} \times \varepsilon_t^m$ as the interaction variable. If the network mechanism operates through supplier linkages rather than some omitted industry characteristic correlated with own upstreamness, this “network neighbor” variable should predict price responses. Panel (b) of [Figure 7](#) is consistent with the supplier-weighted prediction: the interaction is negative and statistically significant during the second year after the shock. The fact that supplier network position—a variable determined entirely by the IO structure of the economy—predicts own-industry price responses is suggestive of network propagation rather than an omitted characteristic of individual sectors.⁴³

⁴³Table A6 and Table A7 of Appendix A-V report the full set of estimates.

Figure 7: Amplification of network effects



Notes: Panel (a) plots the coefficient on the interaction term between upstreamness, input share, and monetary policy shocks. Panel (b) plots the coefficient on the interaction term between supplier-weighted upstreamness and monetary policy shocks. Shaded bands are 90% confidence intervals based on Driscoll-Kraay standard errors.

6 Robustness and extensions

I assess the robustness of the main findings along several dimensions.

Lags, controls, and functional form I start by examining the robustness of the coefficient on the interaction term between upstreamness and monetary policy shocks to using a richer lag structure for the dependent variable and adding lags of the monetary policy shocks. The results are virtually unchanged. Also, I include some typical macro controls (i.e., CPI inflation, industrial production growth, the change in unemployment, and M2 growth), and find similar results. Finally, I also find consistent results when I replace the measure of upstreamness with its log, confirming that the linear specification is not imposing an overly restrictive functional form. [Table A8](#) of [Appendix A-V](#) reports all the estimates.

Measurement I then test the robustness of the results to using other network measures and identification strategies of monetary policy shocks. [Table 2](#) reports results using alternative specifications at the 18-month horizon, which corresponds to the peak of the cumulative price response. Columns (1) through (3) use the Bauer-Swanson shock interacted with alternative network position measures: upstreamness (column 1, baseline), eigenvector centrality (column 2), and the intermediate input share (column 3). Column (4) uses the [Jarociński and Karadi \(2020\)](#) monetary policy shock, which separates pure monetary policy surprises from central bank information effects using sign restrictions on the co-movement of interest rates and stock prices around FOMC announcements, on an extended sample through December 2024. These estimates confirm that

the interaction is not driven by a particular choice of network measure or shock identification.

Table 2: Alternative network measures and monetary policy identification

	Upstreamness (1)	Centrality (2)	Intm. Share (3)	JK MP Shocks (4)
Network measure $\times$ MP shocks	-0.113*** (0.039)	-0.384*** (0.131)	-0.728*** (0.280)	-0.153** (0.059)
Observations	69,772	69,772	69,772	69,772
R^2	0.003	0.003	0.003	0.004

Notes: The dependent variable is the 18-month cumulative price change. Columns (1)–(3) use the Bauer-Swanson shock interacted with alternative network position measures: upstreamness (column 1, baseline), eigenvector centrality (column 2), and intermediate input share (column 3). Column (4) uses the [Jarociński and Karadi \(2020\)](#) monetary policy shocks, which decomposes high-frequency FOMC surprises into pure monetary policy and central bank information components, on an extended sample through 2024. All specifications include industry fixed effects, sector-time fixed effects, and a PPI inflation lag. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$; * $p < 0.10$.

Industry characteristics A natural concern is that upstreamness may proxy for other industry characteristics. For example, labor intensity, energy cost exposure, intermediate input dependence, or international trade exposure can independently generate differential price responses to monetary shocks. [Table 3](#) addresses this through horse-race regressions that include competing interaction terms alongside the upstreamness interaction. Columns (2)–(5) add each confounder individually, and the final column includes all confounders simultaneously. The coefficient on the interaction between upstreamness and monetary policy shocks remains significant across all specifications, indicating that the network position effect is not driven by these industry characteristics. In particular, the inclusion of trade exposure mitigates the concern that the differential price response merely reflects upstream industries’ greater exposure to global commodity markets and exchange rate pass-through.⁴⁴ The energy cost share control is especially relevant for the concern that upstream industries’ prices respond to monetary shocks through commodity-price channels rather than through the network mechanism. In a separate specification that includes only the energy share interaction alongside upstreamness (not reported in the table), the upstreamness coefficient retains roughly 80% of its baseline magnitude at the peak horizon, while the energy share interaction is not significant at any horizon. This mitigates the concern that commodity-price exposure is the primary driver of the cross-sectional pattern.

⁴⁴Trade exposure is measured as the ratio of exports to total output at the six-digit NAICS level, constructed from BEA data. Upstream industries (e.g., primary metals, chemicals) tend to have higher export shares, making this an important control.

Table 3: Including controls for industry characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
Upstreamness $\times$ MP shocks	-0.113*** (0.039)	-0.104*** (0.037)	-0.105*** (0.038)	-0.102*** (0.036)	-0.112*** (0.039)	-0.090*** (0.034)
Labor share $\times$ MP shocks		0.619*** (0.232)				0.468** (0.205)
Energy share $\times$ MP shocks			-0.153 (0.122)			-0.148 (0.120)
Intm. share $\times$ MP shocks				-0.537** (0.227)		-0.361* (0.201)
Trade exposure $\times$ MP shocks					-0.034 (0.094)	-0.041 (0.091)
Observations	69,772	69,772	69,772	69,772	69,772	69,772
R^2	0.003	0.004	0.003	0.004	0.003	0.004

Notes: The dependent variable is the 18-month cumulative price change. Column (1) is the baseline (upstreamness $\times$ MP shock only). Columns (2)–(5) add one confounder at a time: labor share, energy cost share, intermediate input share, and trade exposure, each interacted with the monetary policy shock. Column (6) includes all confounders simultaneously. All specifications include industry fixed effects, sector-time fixed effects, and a PPI inflation lag. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

While the results in Table 3 are reassuring, upstreamness might still proxy for some other industry characteristic that is harder to control for and that generates differential price dynamics even absent monetary shocks. To mitigate this concern, I estimate the interaction term between upstreamness and the monetary policy shocks at negative horizons (starting in $h = -12, \dots, -2$). Under the null hypothesis that upstreamness is pre-determined with respect to the monetary shock, β^h should be zero at negative horizons. Finding significant pre-trends would indicate that the interaction of network position with the monetary shock captures pre-existing differences in price dynamics rather than causal effects.⁴⁵ At all negative horizons, the estimated coefficients are small in magnitude and statistically indistinguishable from zero, with point estimates clustered near zero and confidence bands comfortably covering the null (Table A9 of Appendix A-V). This suggests that the differential price response after the monetary policy shock is triggered by the shock itself rather than by any pre-existing trend in relative prices across upstream and downstream sectors.

Manufacturing subsample and industry outliers The baseline specification pools all 271 industries with available PPI data, including both manufacturing and non-manufacturing sectors. A natural question is whether the upstreamness effect is specific to manufacturing. I repeat the main analysis on the manufacturing subsample (209 industries, NAICS 311–339), finding consistent results (Table A10 of Appendix A-V).

The full cross-section provides substantial protection against the influence of individual indus-

⁴⁵This test is analogous to the pre-trend check in difference-in-differences designs. While the setting here is a continuous treatment (upstreamness) rather than binary, the logic is identical: significant coefficients at negative horizons would cast doubt on the exogeneity of the interaction term.

tries on the results. However, there could be industries that drive the results. For example, one candidate for concern is “Petroleum and Coal Products”, which has high upstreamness and is subject to large oil-price-driven price swings. Thus, I perform a leave-one-out analysis, re-estimating the baseline specification 271 times, each time dropping a different industry. The distribution of leave-one-out estimates one year and a half after the shock is tightly concentrated around the baseline point estimate (Figure A3 of Appendix A-V), confirming that no single industry drives the results.

Placebo test To ensure that the results are not an artifact of the empirical methodology, I randomly permute the upstreamness rankings across industries 200 times and re-estimate Equation 27. At the horizons where the main results are significant (6–24 months), the actual estimate lies far in the tail of the permutation distribution, well beyond the 5th and 95th percentile bounds (Figure A4 of Appendix A-V). This confirms that it is the specific pattern of upstreamness in the real IO table—not an arbitrary industry ranking—that generates the differential price response.

Sub-samples I also verify that the results are not driven by any particular episode. Table A11 of Appendix A-V reports the estimates a year and a half after the shock across different sub-samples. The pre-crisis sub-sample (through 2010) yields a coefficient of -0.073 , the post-crisis sub-sample (2011 onward) yields -0.118 , and the sub-sample excluding the COVID period (through December 2019) yields -0.104 , all statistically significant. The stability of the estimates across sub-periods supports the interpretation that the network amplification mechanism is a structural feature of the economy.

Demand cyclicalities A concern is that upstream industries within a 3-digit sector may have inherently more volatile demand—for example, because their customers in capital goods are more interest-rate sensitive—and this demand cyclicalities, rather than the network mechanism, drives the interaction coefficient. To address this, I compute each industry’s PPI volatility in the early sample (1998–2002) and include its interaction with the monetary shock as an additional control. The sample drops to 42,000 observations because some industries entered BLS PPI coverage after 2002; in this restricted sample, the upstreamness interaction retains its sign, though it falls to roughly half its baseline size (-0.055 at $h = 18$, compared to -0.113 in the baseline), and the PPI volatility control itself is not significant at any horizon. The reduction in both magnitude and precision reflects the much smaller sample; the insignificance of the volatility control indicates that demand cyclicalities does not account for the upstreamness effect.

7 Conclusion

This paper shows that the production network is a first-order determinant of how monetary policy transmits to prices. More upstream industries—those farther from final demand in the IO

tables—exhibit significantly larger cumulative price responses to identified monetary shocks. A multi-sector New Keynesian model with IO linkages rationalizes this finding by showing that the network position determines each sector’s customer composition (B2B versus B2C), which in turn determines its price flexibility. A counterfactual decomposition shows that this network-determined rigidity heterogeneity, rather than cost-cascade propagation through IO linkages, is the primary driver of the cross-sectional pattern.

The paper also shows that this cross-sectional pattern operates asymmetrically. Cumulated over the one-to-two-year window, the upstreamness differential following expansionary shocks is large and negative while the contractionary differential is near zero. This asymmetry, driven by asymmetric price rigidity interacting with the network-determined rigidity differential, provides micro-founded, cross-sectional evidence for why inflation accelerates more easily than it decelerates. For central banks, the production network creates a structural asymmetry in the transmission mechanism, whereby rate cuts generate larger cross-sectional price dispersion than rate hikes.

Because the distribution of price rigidity is endogenous to network structure, average rigidity is no longer a sufficient statistic for monetary potency, and tracking how the network evolves over time becomes part of the inflation-monitoring problem. In this regard, two long-running shifts are particularly relevant. As the services share of GDP continues to rise, the economy moves toward sectors that are both more downstream and stickier, widening the upstream-downstream wedge that drives the patterns documented here. As global value chains lengthen or re-shore, the number of intermediate links between producers and consumers changes, and with it the room for asymmetric rigidity to compound stage by stage.

References

- ACEMOGLU, D., CARVALHO, V. M., OZDAGLAR, A. and TAHBAZ-SALEHI, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, **80** (5), 1977–2016.
- AFROUZI, H. and BHATTARAI, S. (2024). Inflation and GDP dynamics in production networks: A sufficient statistics approach. *NBER Working Paper No. 31218*.
- ANTRÀS, P., CHOR, D., FALLY, T. and HILLBERRY, R. (2012). Measuring the upstreamness of production and trade flows. *American Economic Review*, **102** (3), 412–416.
- ATALAY, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, **9** (4), 254–280.
- BAQAEE, D. R. and FARHI, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten’s theorem. *Econometrica*, **87** (4), 1155–1203.
- and — (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics*, **135** (1), 105–163.
- BARROT, J.-N. and SAUVAGNAT, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *The Quarterly Journal of Economics*, **131** (3), 1543–1592.
- BASU, S. (1995). Intermediate goods and business cycles: Implications for productivity and welfare. *American Economic Review*, **85** (3), 512–531.
- BAUER, M. D. and SWANSON, E. T. (2023). An alternative explanation for the “Fed information effect”. *American Economic Review*, **113** (3), 664–700.
- BILS, M. and KLENOW, P. J. (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy*, **112** (5), 947–985.
- CABALLERO, R. J. and ENGEL, E. M. R. A. (1999). Explaining investment dynamics in U.S. manufacturing: A generalized (S,s) approach. *Econometrica*, **67** (4), 783–826.
- CALVO, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, **12** (3), 383–398.
- CARVALHO, C. (2006). Heterogeneity in price stickiness and the real effects of monetary shocks. *The B.E. Journal of Macroeconomics*, **6** (1), 1–58.
- CARVALHO, V. M. (2014). From micro to macro via production networks. *Journal of Economic Perspectives*, **28** (4), 23–48.
- , NIREI, M., SAITO, Y. U. and TAHBAZ-SALEHI, A. (2021). Supply chain disruptions: Evidence from the great East Japan earthquake. *The Quarterly Journal of Economics*, **136** (2), 1255–1321.
- and SCHWARTZMAN, F. (2015). Selection and monetary non-neutrality in time-dependent pricing models. *Journal of Monetary Economics*, **76**, 141–156.
- and TAHBAZ-SALEHI, A. (2019). Production networks: A primer. *Annual Review of Economics*, **11**, 635–663.
- CAVALLO, A., LIPPI, F. and MIYAHARA, K. (2024). Large and small price changes. *American Economic Review*, forthcoming.
- CHRISTIANO, L. J., EICHENBAUM, M. and EVANS, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, **113** (1), 1–45.

- GABAIX, X. (2011). The granular origins of aggregate fluctuations. *Econometrica*, **79** (3), 733–772.
- GALÍ, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton, NJ: Princeton University Press, 2nd edn.
- GHAASSIBE, M. (2023). Monetary policy and production networks: An empirical investigation. *Journal of Monetary Economics*, **133**, 140–159.
- (2025). *Endogenous Production Networks and Non-Linear Monetary Transmission*. Tech. rep., University of Oxford, working paper.
- GOLOSOV, M. and LUCAS, R. E., JR. (2007). Menu costs and Phillips curves. *Journal of Political Economy*, **115** (2), 171–199.
- GOPINATH, G. and ITSKHOKI, O. (2010). Frequency of price adjustment and pass-through. *The Quarterly Journal of Economics*, **125** (2), 675–727.
- JAROCIŃSKI, M. and KARADI, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, **12** (2), 1–43.
- JORDÀ, O. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, **95** (1), 161–182.
- KEHOE, P. and MIDRIGAN, V. (2015). Prices are sticky after all. *Journal of Monetary Economics*, **75**, 35–53.
- LA’O, J. and TAHBAZ-SALEHI, A. (2022). Optimal monetary policy in production networks. *Econometrica*, **90** (3), 1295–1336.
- LONG, J. B., JR. and PLOSSER, C. I. (1983). Real business cycles. *Journal of Political Economy*, **91** (1), 39–69.
- MANKIW, N. G. and REIS, R. (2002). Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, **117** (4), 1295–1328.
- MIDRIGAN, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, **79** (4), 1139–1180.
- MINTON, R. and WHEATON, B. (2023). Delayed inflation in supply chains: Theory and evidence. *Working Paper*.
- NAKAMURA, E. and STEINSSON, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, **123** (4), 1415–1464.
- and — (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics*, **125** (3), 961–1013.
- and — (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, **133** (3), 1283–1330.
- OZDAGLI, A. K. and WEBER, M. (2025). Monetary policy through production networks: Evidence from the stock market. *The Review of Financial Studies*, advance article.
- PASTEN, E., SCHOENLE, R. and WEBER, M. (2020). The propagation of monetary policy shocks in a heterogeneous production economy. *Journal of Monetary Economics*, **116**, 1–22.
- PASTÉN, E., SCHOENLE, R. and WEBER, M. (2024). Sectoral heterogeneity in nominal price rigidity and the origin of aggregate fluctuations. *American Economic Journal: Macroeconomics*, **16** (2), 1–41.

- PELTZMAN, S. (2000). Prices rise faster than they fall. *Journal of Political Economy*, **108** (3), 466–502.
- PLAGBORG-MØLLER, M. and WOLF, C. K. (2021). Local projections and VARs estimate the same impulse responses. *Econometrica*, **89** (2), 955–980.
- ROMER, C. D. and ROMER, D. H. (2004). A new measure of monetary shocks: Derivation and implications. *American Economic Review*, **94** (4), 1055–1084.
- RUBBO, E. (2023). Networks, Phillips curves, and monetary policy. *Econometrica*, **91** (4), 1417–1455.
- TENREYRO, S. and THWAITES, G. (2016). Pushing on a string: US monetary policy is less powerful in recessions. *American Economic Journal: Macroeconomics*, **8** (4), 43–74.
- WOODFORD, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.

Online Appendix for:
“Production Networks and the
(Asymmetric) Transmission of Monetary Policy”

Francesco Grigoli

A-I Proofs

This appendix provides formal proofs of the propositions stated in the main text.

A-I.1 Proof of [Proposition 1](#): sectoral Phillips curves

Consider sector i with Calvo parameter θ_i . In each period, a fraction $1 - \theta_i$ of firms reset their price. A resetting firm j chooses p_{ijt}^* to maximize

$$\mathbb{E}_t \sum_{s=0}^{\infty} (\beta\theta_i)^s \left[p_{ijt}^* y_{ij,t+s} - \text{TC}_{ij,t+s}(y_{ij,t+s}) \right]. \quad (34)$$

The first-order condition equates the expected present value of the markup to the desired markup. Log-linearizing around the zero-inflation steady state:

$$\hat{p}_{ijt}^* = (1 - \beta\theta_i) \sum_{s=0}^{\infty} (\beta\theta_i)^s \mathbb{E}_t [\widehat{mc}_{i,t+s} + \hat{p}_{i,t+s}], \quad (35)$$

where $\widehat{mc}_{i,t+s}$ is the log deviation of real marginal cost.

The Cobb-Douglas production function in [Equation 3](#), evaluated at the zero-inflation steady state with $\hat{z}_{it} = 0$ (no sectoral productivity shocks), implies that the log-linearized nominal marginal cost is

$$\widehat{mc}_{it} = \alpha_i \hat{w}_t + \sum_{k=1}^N \gamma_{ik} \hat{p}_{kt}, \quad (36)$$

so the real marginal cost relative to sector i 's own price is $\widehat{mc}_{it} = \alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt} - \hat{p}_{it}$.

The CES price index for sector i evolves as $\hat{p}_{it} = \theta_i \hat{p}_{i,t-1} + (1 - \theta_i) \hat{p}_{it}^*$. Combining the optimal reset price, the marginal cost expression, and the price index law of motion yields the sectoral Phillips curve in [Equation 6](#) with slope $\lambda_i = (1 - \theta_i)(1 - \beta\theta_i)/\theta_i$. $\square$

A-I.2 Proof of [Proposition 2](#): network-adjusted rigidity

The aggregate gross-output price index is defined as the Domar-weighted sum of sectoral prices: $\hat{P}_t = \sum_{i=1}^N \xi_i \hat{p}_{it}$, where the Domar weight $\xi_i = Y_i/\text{GDP}$ is the ratio of sector i 's gross output to GDP.

Differentiating with respect to the monetary shock ε_t^m :

$$\frac{\partial \hat{p}_{t+h}}{\partial \varepsilon_t^m} = \sum_{i=1}^N \zeta_i \cdot \frac{\partial \hat{p}_{i,t+h}}{\partial \varepsilon_t^m}, \quad (37)$$

which is Equation 9. The output effect follows from the quantity theory in log-deviations: $\hat{Y}_{t+h} = \hat{D}_{t+h} - \hat{p}_{t+h}$, where $\hat{D}_{t+h}$ captures the nominal demand impulse. Output non-neutrality ($\hat{Y}_{t+h} \neq 0$) requires $\hat{p}_{t+h} \neq \hat{D}_{t+h}$, which obtains when the Domar-weighted price level adjusts slowly.

The alignment result (Corollary 1) follows from the observation that sectors with higher θ_i (more rigid prices) have smaller λ_i (flatter Phillips curves) and hence slower price adjustment. When these high- θ_i sectors also have high Domar weights ζ_i , the Domar-weighted aggregate price $\hat{p}_{t+h}$ adjusts slowly, maximizing the gap $\hat{D}_{t+h} - \hat{p}_{t+h}$ and hence aggregate output non-neutrality. $\square$

A-I.3 Proof of Proposition 3: upstreamness and price dynamics

The proof proceeds under Assumption 1, stated in the main text: $r_i \equiv \lambda_i / (1 + \lambda_i)$ is weakly increasing in U_i .

Part (i): general-network factorization. In the static case ($\beta = 0$), the sectoral Phillips curve for sector i is

$$\hat{p}_{it} = \lambda_i \left[\alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt} - \hat{p}_{it} \right]. \quad (38)$$

Rearranging: $(1 + \lambda_i) \hat{p}_{it} = \lambda_i [\alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt}]$, so

$$\hat{p}_{it} = r_i \underbrace{\left[\alpha_i \hat{w}_t + \sum_{k=1}^N \gamma_{ik} \hat{p}_{kt} \right]}_{\equiv \widehat{nm}c_{it}}, \quad (39)$$

where $r_i = \lambda_i / (1 + \lambda_i)$ and $\widehat{nm}c_{it}$ is sector i 's nominal marginal cost deviation.

In matrix form, defining $\mathbf{R} = \text{diag}(r_1, \dots, r_N)$ and $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_N)'$:

$$\hat{\mathbf{p}}_t = (\mathbf{I} - \mathbf{R}\boldsymbol{\Gamma})^{-1} \mathbf{R}\boldsymbol{\alpha} \hat{w}_t. \quad (40)$$

The matrix $(\mathbf{I} - \mathbf{R}\boldsymbol{\Gamma})^{-1}$ exists and has non-negative entries because $\mathbf{R}\boldsymbol{\Gamma}$ is a non-negative matrix with spectral radius strictly less than one (since $r_i < 1$ and $\sum_k \gamma_{ik} < 1$ for all i). It follows that $\hat{p}_{it} / \hat{w}_t > 0$ for all i : every sector's price moves in the same direction as the wage. Hence $\widehat{nm}c_{it} = \alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt}$ has the same sign as $\hat{w}_t$ and is strictly nonzero.

The factorization in Equation 39 shows that sector i 's price response is the product of two positive terms: its own flexibility r_i and its marginal cost response $|\widehat{nm}c_{it}|$. Under Assumption 1, r_i is

increasing in U_i , so the flexibility channel generates a positive cross-sectional association between upstreamness and $|\hat{p}_{it}|$.

The input-cost structure creates a second channel operating through $\widehat{nm\bar{c}}_{it}$. Since $\hat{p}_{kt}$ enters sector i 's marginal cost with weight γ_{ik} , the marginal cost response varies across sectors with their IO linkages. The sign of this channel's contribution to the upstreamness differential depends on the joint distribution of labor shares and network position. In the stylized chain of Part (ii), where the most upstream sector is the most labor-intensive ($\alpha_1 = 1$), it reinforces the flexibility channel. In U.S. data, where upstream sectors have lower labor shares ($\text{Corr}(U_i, \alpha_i) = -0.90$), their marginal cost depends more on other sticky prices and less on the flexible wage, so this channel partially offsets the flexibility channel, as the decomposition in [Subsection 5.2](#) shows. In both cases the flexibility channel is the quantitatively dominant force.

Part (ii): chain-economy closed form. Consider a chain of K sectors where sector 1 uses only labor ($\alpha_1 = 1$) and sector $k \geq 2$ uses sector $k-1$'s output with cost share γ ($\alpha_k = 1 - \gamma$, $\gamma_{k,k-1} = \gamma$). Assume uniform flexibility r across sectors. Note that this chain assigns the most upstream sector the highest labor share, the reverse of the U.S. pattern; it isolates the cost-compounding geometry by shutting off the flexibility channel.

Define $x_k \equiv \hat{p}_{kt}/\hat{w}_t$. From [Equation 39](#):

$$x_1 = r, \quad x_k = r[(1 - \gamma) + \gamma x_{k-1}] \quad \text{for } k \geq 2. \quad (41)$$

This is a first-order linear recursion $x_k = \phi x_{k-1} + r(1 - \gamma)$ with $\phi \equiv r\gamma \in (0, 1)$. The solution is

$$x_k = \frac{r[(1 - \gamma) + \gamma(1 - r)\phi^{k-1}]}{1 - \phi}. \quad (42)$$

Verification: $x_1 = r[(1 - \gamma) + \gamma(1 - r)]/(1 - \phi) = r(1 - r\gamma)/(1 - r\gamma) = r$. ✓

Since ϕ^{k-1} is strictly decreasing in k and enters with a positive coefficient $\gamma(1 - r) > 0$, the sequence x_k is strictly decreasing: more downstream sectors have smaller price responses.

The upstream-downstream differential is

$$D_K \equiv x_1 - x_K = \frac{r\gamma(1 - r)[1 - \phi^{K-1}]}{1 - \phi}. \quad (43)$$

Verification: $D_2 = r\gamma(1 - r)(1 - \phi)/(1 - \phi) = r\gamma(1 - r)$. With $r = 0.3$, $\gamma = 0.5$: $D_2 = 0.3 \times 0.5 \times 0.7 = 0.105 = 0.3 - 0.195$. ✓

Since ϕ^{K-1} is decreasing in K , D_K is increasing in K and converges to $r\gamma(1 - r)/(1 - \phi) < \infty$.

Part (iii): delayed peak output effect in a chain economy. The real output response of sector k satisfies $\hat{y}_{kt} \approx \hat{d}_{kt} - \hat{p}_{kt}$, where $\hat{d}_{kt}$ is the demand impulse. From Part (ii), $|\hat{p}_{kt}|$ is decreasing in k : downstream sectors have smaller price responses. Holding the demand impulse constant

across sectors, a smaller $|\hat{p}_{kt}|$ means that the output gap $\hat{d}_{kt} - \hat{p}_{kt}$ is larger and persists longer for downstream sectors. The peak real effect therefore occurs at a later horizon for more downstream sectors.

Remark on general networks. Parts (ii) and (iii) are proven for chain economies. In the calibrated model using the full BEA IO matrix (Subsection 5.2), I verify numerically that the cross-sectional price response is strongly positively correlated with upstreamness and that the peak output timing is later for downstream sectors, confirming that the chain-economy predictions extend to the empirically relevant network structure. $\square$

A-I.4 Proof of Proposition 5: sign-dependent network propagation

In the static case ($\beta = 0$), the sectoral Phillips curve $\hat{p}_{it} = \lambda_i^s[\alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt} - \hat{p}_{it}]$ rearranges to $\hat{p}_{it} = r_i^s[\alpha_i \hat{w}_t + \sum_k \gamma_{ik} \hat{p}_{kt}]$ with $r_i^s = \lambda_i^s / (1 + \lambda_i^s)$. In matrix form, $(\mathbf{I} - \mathbf{R}^s \Gamma) \hat{\mathbf{p}}^s = \mathbf{R}^s \alpha \hat{\Delta} w$, which gives Equation 23 provided $(\mathbf{I} - \mathbf{R}^s \Gamma)$ is invertible.

Existence and the Neumann series. The matrix $\mathbf{R}^s \Gamma$ is non-negative. Its row sums are $r_i^s \sum_k \gamma_{ik} = r_i^s (1 - \alpha_i) < 1$, since $r_i^s < 1$ and $\alpha_i > 0$. A non-negative matrix with all row sums below one has spectral radius below one, so $(\mathbf{I} - \mathbf{R}^s \Gamma)^{-1} = \sum_{k=0}^{\infty} (\mathbf{R}^s \Gamma)^k$ exists with non-negative entries, and

$$\mathbf{M}^s = (\mathbf{I} - \mathbf{R}^s \Gamma)^{-1} \mathbf{R}^s = \sum_{k=0}^{\infty} (\mathbf{R}^s \Gamma)^k \mathbf{R}^s. \quad (44)$$

Part (i): nesting. If $r_i^+ = r_i^-$ for all i , then $\mathbf{R}^+ = \mathbf{R}^-$ and hence $\mathbf{M}^+ = \mathbf{M}^-$ termwise in Equation 44. Price responses are then identical across shock signs.

Part (ii): elementwise monotonicity. Each term $(\mathbf{R}^s \Gamma)^k \mathbf{R}^s$ in Equation 44 is a product of the non-negative diagonal matrix $\mathbf{R}^s$ (appearing $k+1$ times) and the non-negative matrix Γ (k times). Every entry of such a product is a sum of monomials in $\{r_i^s\}$ and $\{\gamma_{ik}\}$ with non-negative coefficients, and is therefore non-decreasing in each r_i^s , holding Γ fixed. Consequently, if $r_i^+ \geq r_i^-$ for all i , then every term satisfies $(\mathbf{R}^+ \Gamma)^k \mathbf{R}^+ \geq (\mathbf{R}^- \Gamma)^k \mathbf{R}^-$ elementwise, and summing over k gives $\mathbf{M}^+ \geq \mathbf{M}^-$ elementwise. Since $\alpha \geq \mathbf{0}$, it follows that $|\hat{p}_i^+| = (\mathbf{M}^+ \alpha)_i |\hat{\Delta} w| \geq (\mathbf{M}^- \alpha)_i |\hat{\Delta} w| = |\hat{p}_i^-|$ for every sector i . $\square$

Remark. Elementwise dominance $\mathbf{M}^+ \geq \mathbf{M}^-$ implies a uniformly larger response but does not by itself order the cross-sectional regression slopes of $\hat{p}_i^s$ on U_i . The chain economy below shows the slope (upstream-downstream differential) is steeper under expansionary shocks in closed form, and Subsection 5.2 verifies numerically that the same ordering holds for the calibrated BEA network.

A-I.5 Proof of Proposition 6: asymmetric network amplification

Consider a chain of K sectors indexed $1, \dots, K$, where sector k uses only sector $k-1$'s output (and labor) as inputs, and sector 1 uses only labor. The IO matrix has $\gamma_{k,k-1} = \gamma$ for $k \geq 2$ and all other entries zero.

Under asymmetric Calvo pricing with $\beta = 0$, the state-dependent price-adjustment probability is $r^s = \lambda^s / (1 + \lambda^s)$ and the cascade factor is $\phi^s = r^s \gamma$, where $s \in \{+, -\}$.

Part (i): derivation of D_K^s . Define $x_k^s \equiv \hat{p}_k^s / \Delta w$. Sector 1 (pure labor): $x_1^s = r^s$. For $k \geq 2$: $x_k^s = r^s[(1 - \gamma) + \gamma x_{k-1}^s]$. This is the same recursion as in the proof of [Proposition 3\(ii\)](#), with r replaced by r^s . The solution is

$$x_k^s = \frac{r^s [(1 - \gamma) + \gamma(1 - r^s)(\phi^s)^{k-1}]}{1 - \phi^s}. \quad (45)$$

Verification: $x_1^s = r^s(1 - \phi^s) / (1 - \phi^s) = r^s$. ✓

The upstream-downstream differential is

$$D_K^s \equiv x_1^s - x_K^s = \frac{r^s \gamma (1 - r^s) [1 - (\phi^s)^{K-1}]}{1 - \phi^s}. \quad (46)$$

Verification ($K = 2, r^+ = 0.4, r^- = 0.2, \gamma = 0.5$): $D_2^+ = 0.4 \times 0.5 \times 0.6 = 0.12$; $D_2^- = 0.2 \times 0.5 \times 0.8 = 0.08$. Direct calculation: $x_1^+ = 0.4, x_2^+ = 0.4(0.5 + 0.5 \times 0.4) = 0.28, D_2^+ = 0.12$. ✓

Part (ii): the asymmetry ratio $\mathcal{R}_K$. Using the identity $[1 - (\phi^s)^{K-1}] / (1 - \phi^s) = \sum_{j=0}^{K-2} (\phi^s)^j$, the differential simplifies to

$$D_K^s = r^s \gamma (1 - r^s) \sum_{j=0}^{K-2} (\phi^s)^j. \quad (47)$$

The ratio of the expansionary to the contractionary differential is therefore

$$\mathcal{R}_K = \frac{D_K^+}{D_K^-} = \underbrace{\frac{r^+(1 - r^+)}{r^-(1 - r^-)}}_{\text{own-flexibility factor}} \times \underbrace{\frac{\sum_{j=0}^{K-2} (\phi^+)^j}{\sum_{j=0}^{K-2} (\phi^-)^j}}_{\equiv \mathcal{G}(K) : \text{network factor}}. \quad (48)$$

The γ terms cancel exactly, so $\mathcal{R}_K$ is independent of the input share.

Verification ($r^+ = 0.4, r^- = 0.2, \gamma = 0.5$): $\mathcal{R}_2 = [0.24/0.16] \times 1 = 1.5$; $\mathcal{R}_3 = 1.5 \times (1.2/1.1) = 1.636$. Direct: $D_3^+ / D_3^- = 0.144/0.088 = 1.636$. ✓

Properties of $\mathcal{R}_K$ under $r^+, r^- < \frac{1}{2}$ with $r^+ > r^-$. (a) $\mathcal{R}_K > 1$. The function $f(r) = r(1 - r)$ is strictly increasing on $[0, \frac{1}{2}]$, so the own-flexibility factor exceeds one. At $K = 2$, $\mathcal{G}(2) = 1$, so $\mathcal{R}_2 = r^+(1 - r^+) / [r^-(1 - r^-)] > 1$. For $K > 2$, $\mathcal{G}(K) \geq 1$ (shown below), so $\mathcal{R}_K \geq \mathcal{R}_2 > 1$.

(b) $\mathcal{G}(K)$ is strictly increasing in K . Define $S_n^s = \sum_{j=0}^n (\phi^s)^j$. Then $\mathcal{G}(K) = S_{K-2}^+ / S_{K-2}^-$. The claim $\mathcal{G}(K+1) > \mathcal{G}(K)$ holds if and only if $(\phi^+ / \phi^-)^{K-1} > \mathcal{G}(K)$. Now $\mathcal{G}(K)$ is a weighted average of $(\phi^+ / \phi^-)^j$ for $j = 0, \dots, K-2$, with weights $(\phi^-)^j / S_{K-2}^-$ that are decreasing in j (since $\phi^- < 1$).

A weighted average with decreasing weights is bounded above by its largest term, $(\phi^+/\phi^-)^{K-2}$. Since $\phi^+/\phi^- > 1$, it follows that $(\phi^+/\phi^-)^{K-1} > (\phi^+/\phi^-)^{K-2} \geq \mathcal{G}(K)$, completing the argument.

(c) *Finite limit.* As $K \rightarrow \infty$, $S_{K-2}^s \rightarrow 1/(1 - \phi^s)$, so $\mathcal{G}(\infty) = (1 - \phi^-)/(1 - \phi^+)$ and $\mathcal{R}_\infty = r^+(1 - r^+)(1 - \phi^-)/[r^-(1 - r^-)(1 - \phi^+)]$. $\square$

A-II The network Phillips curve

Return to the matrix system of sectoral Phillips curves from [Proposition 1](#):

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \Lambda [\mathcal{A} \hat{w}_t \mathbf{1} + (\Gamma - \mathbf{I}) \hat{\mathbf{p}}_t]. \quad (49)$$

To derive the aggregate Phillips curve, define aggregate inflation as the Domar-weighted average $\pi_t^{\text{agg}} = \xi' \pi_t$, where ξ is the vector of Domar weights from [Corollary 1](#).

The derivation of the aggregate slope proceeds in three steps. First, note that the sectoral price vector and sectoral inflation are linked by $\pi_t = \hat{\mathbf{p}}_t - \hat{\mathbf{p}}_{t-1}$. In the static version of the model (setting $\beta = 0$ for clarity; the full dynamic result follows by identical algebra with $\beta > 0$), [Equation 49](#) reduces to

$$\pi_t = \Lambda [\mathcal{A} \hat{w}_t \mathbf{1} + (\Gamma - \mathbf{I}) \hat{\mathbf{p}}_t]. \quad (50)$$

Second, since $\hat{\mathbf{p}}_t$ is a function of the history of sectoral inflation rates, consider the impact response to a nominal demand impulse $\hat{w}_t = x_t$ (where x_t denotes the aggregate output gap). Solving [Equation 50](#) for the price vector yields

$$\hat{\mathbf{p}}_t = [\mathbf{I} + \Lambda(\mathbf{I} - \Gamma)]^{-1} \Lambda \mathcal{A} x_t \mathbf{1} = [\mathbf{I} - (\mathbf{I} - \Lambda)\Gamma]^{-1} \Lambda \alpha x_t, \quad (51)$$

where $\alpha = (\alpha_1, \dots, \alpha_N)'$ is the vector of labor shares, and the second equality uses the matrix identity $[\mathbf{I} + \Lambda(\mathbf{I} - \Gamma)]^{-1} \Lambda \mathcal{A} = [\mathbf{I} - (\mathbf{I} - \Lambda)\Gamma]^{-1} \Lambda \mathcal{A}$, which holds because both expressions invert the same mapping from marginal costs to prices.⁴⁶

Third, aggregate using Domar weights. Since $\pi_t^{\text{agg}} = \xi' \pi_t$ and each π_{it} equals $\hat{p}_{it}$ in the impact period, the aggregate Phillips curve takes the form

$$\pi_t^{\text{agg}} = \beta \mathbb{E}_t \pi_{t+1}^{\text{agg}} + \kappa_{\text{net}} x_t, \quad (52)$$

where the *network slope* is

$$\kappa_{\text{net}} = \xi' [\mathbf{I} - (\mathbf{I} - \Lambda)\Gamma]^{-1} \Lambda \alpha. \quad (53)$$

The matrix $(\mathbf{I} - \Lambda)\Gamma$ captures the interaction between nominal rigidity and the IO structure. Its (i, k) entry is $(1 - \lambda_i) \gamma_{ik}$: sector i 's reliance on sector k 's output, discounted by sector i 's rigidity. The Neumann series expansion of the inverse, $[\mathbf{I} - (\mathbf{I} - \Lambda)\Gamma]^{-1} = \sum_{n=0}^{\infty} [(\mathbf{I} - \Lambda)\Gamma]^n$, shows that κ_{net} sums the cost pass-through across all production chain lengths, with rigidity attenuating the pass-through at each stage.

Proposition 7 (Network Phillips Curve). *The aggregate Phillips curve [Equation 52](#) has slope κ_{net} given*

⁴⁶To see this, write $\mathbf{I} + \Lambda(\mathbf{I} - \Gamma) = \Lambda[\Lambda^{-1} + \mathbf{I} - \Gamma] = \Lambda[\Lambda^{-1}(\mathbf{I} - (\mathbf{I} - \Lambda)\Gamma)]$. Taking inverses and right-multiplying by $\Lambda \mathcal{A}$ yields the result, noting that $\mathcal{A} = \text{diag}(\alpha)$.

by Equation 53. This slope satisfies:

- (i) $\kappa_{\text{net}} = \kappa_{\text{standard}}$ when $\Gamma = \mathbf{0}$: with no intermediate input linkages, the network Phillips curve reduces to the standard single-sector New Keynesian Phillips curve with slope $\kappa_{\text{standard}} = \zeta' \Lambda \alpha$.
- (ii) $\kappa_{\text{net}} < \kappa_{\text{standard}}$ when Domar-weighted upstream sectors are more rigid (i.e., when sectors with high ζ_i and high U_i have high θ_i): the production network flattens the aggregate Phillips curve.
- (iii) $\partial \kappa_{\text{net}} / \partial \rho(\Gamma) < 0$: a denser IO network (higher spectral radius of Γ) reduces the slope of the aggregate Phillips curve.

The formal proof iterates the sectoral Phillips curves through the Neumann series and aggregates using Domar weights; the argument parallels the derivation above with $\beta > 0$. Part (i) is immediate: when $\Gamma = \mathbf{0}$, the inverse in Equation 53 reduces to the identity and $\kappa_{\text{net}} = \zeta' \Lambda \alpha$, which is the weighted-average sectoral slope. For part (ii), observe that inserting a positive Γ into $[\mathbf{I} - (\mathbf{I} - \Lambda)\Gamma]^{-1}$ amplifies the contribution of rigid sectors (high $1 - \lambda_i$) that sit upstream. When these rigid sectors also carry large Domar weights, the inverse tilts κ_{net} below its no-network benchmark. Part (iii) follows because the spectral radius $\rho(\Gamma)$ governs the rate at which the Neumann series decays: a higher $\rho(\Gamma)$ means that longer production chains contribute more to the sum, and each additional stage applies the rigidity discount $(1 - \lambda_i)$, further attenuating the slope.

Remark 1 (Policy implication). *Proposition 7 shows that the production network does not merely redistribute the price response across sectors—it changes the aggregate trade-off between inflation and output. A longer, more interconnected supply chain flattens the aggregate Phillips curve, meaning that monetary policy must move interest rates by more to achieve a given change in inflation. This result complements Rubbo (2023), who derives the aggregate Phillips curve in a multi-sector model with strategic complementarities. The contribution here is to provide the explicit dependence on the spectral radius of the rigidity-weighted IO matrix and to show that the slope is a decreasing function of network density. While Rubbo’s framework emphasizes the role of markup complementarities, Equation 53 isolates the purely mechanical channel through which the IO structure attenuates cost pass-through across production stages, compounding the effect of sectoral rigidities into a flatter aggregate trade-off.*

A-III Menu cost model with production networks

This appendix develops a multi-sector menu cost model with IO linkages as an alternative to the Calvo framework in the main text, building on [Nakamura and Steinsson \(2010\)](#) and [Goloso and Lucas \(2007\)](#). The model replaces random Calvo adjustment with an Ss pricing rule and adds strategic complementarities in price setting. The goal is to assess whether the network amplification mechanism is robust to state-dependent pricing, and to identify additional channels—absent from Calvo models—through which the production network amplifies monetary non-neutrality.

A-III.1 Environment and pricing

The production side is identical to [Subsection 2.1](#): N sectors connected by an IO matrix $\Gamma = [\gamma_{ik}]$, each containing a continuum of monopolistically competitive firms. Log nominal marginal cost for firm j in sector i is $mc_{ijt} = \alpha_i w_t + \sum_k \gamma_{ik} P_{kt} - z_{ijt}$, where z_{ijt} follows a random walk with volatility σ_z .

Each firm's desired price incorporates strategic complementarities following [Gopinath and Itskhoki \(2010\)](#):

$$p_{ijt}^* = (1 - \psi) mc_{ijt} + \psi P_{it}, \quad (54)$$

where $\psi \in [0, 1)$ is the strategic complementarity parameter and P_{it} is the sectoral price index. The price gap $g_{ijt} \equiv p_{ijt}^* - p_{ijt}$ evolves as

$$\Delta g_{ijt} = (1 - \psi) [\alpha_i \Delta w_t + IO_{it} + \mu - \Delta z_{ijt}] + \psi \Delta P_{it}, \quad (55)$$

where $\mu > 0$ is trend inflation and $IO_{it} = \sum_k \gamma_{ik} \Delta P_{kt}$ is the IO-weighted upstream cost signal.

Firms face a fixed menu cost F_i of adjusting their nominal price. With trend inflation, the optimal policy is an asymmetric Ss rule: the firm adjusts when g_{ijt} exceeds an upper threshold $y_i^* + b_i^u$ or falls below $y_i^* - b_i^\ell$, where $y_i^* < 0$ is the return point and $b_i^u > b_i^\ell$. After adjustment, the gap resets to y_i^* .

A-III.2 Calibration

[Table A1](#) reports the parameter values. The Ss bands are calibrated to match the Calvo-equivalent adjustment frequencies from the main text ($\theta_i = \bar{\theta} + \delta(U_i - \bar{U})$, with $\bar{\theta} = 0.85$ and $\delta = -0.152$ from [Pasten et al. \(2020\)](#)). The IO matrix is the BEA Use Table aggregated to $N = 36$ summary sectors.

Impulse responses are computed using the MIT shock approach: the economy is simulated for 2,000 periods to reach the ergodic distribution, then shocked and baseline paths (with identical idiosyncratic draws) are compared over 36 months. Interaction coefficients are obtained from cross-sectional regressions of sectoral IRFs on upstreamness, exactly mirroring the empirical local

Table A1: Menu cost model parameters


Parameter	Value	Description	Source
σ_z	0.04	Idiosyncratic volatility (monthly)	Golosov and Lucas (2007)
μ	0.002	Trend inflation (monthly, $\approx 2.4\%$ annual)	Sample mean PPI inflation
demand path	hump, trough $h=17$	Nominal demand impulse	Estimated IP response to shock
ψ	0.55	Strategic complementarity	Gopinath and Itskhoki (2010)
$\bar{\theta}$	0.85	Avg. Calvo parameter (for Ss calibration)	Pasten <i>et al.</i> (2020) frequencies
δ	−0.152	Upstreamness–flexibility slope	Pasten <i>et al.</i> (2020) frequencies

projection.

A-III.3 Results

[Figure A1](#) compares the interaction coefficients from the menu cost model, the Calvo model, and the empirical estimates.

Figure A1: Menu cost model vs. Calvo: interaction coefficients



Notes: Panel (a) plots the interaction coefficient $\hat{\beta}_h$ from the menu cost model with strategic complementarities ($\psi = 0.55$), the Calvo model (from [Subsection 5.2](#)), and the empirical estimates with 90% confidence band. Panel (b) shows the average fraction of firms adjusting prices at each horizon for upstream (top quartile) and downstream (bottom quartile) sectors, illustrating the endogenous rigidity dispersion generated by the Ss pricing rule.

The menu cost model generates a peak interaction coefficient of -0.125 at $h = 15$, compared to -0.105 for the externally-calibrated Calvo model and -0.113 in the data. Two mechanisms contribute to the slightly larger response. First, endogenous rigidity dispersion: even with identical menu costs, upstream sectors adjust more frequently because they face larger and more volatile cost shocks—driven by their position further from final demand and fewer downstream customers to average over. This generates the cross-sectional variation in rigidity that must be

imposed exogenously through the parameter δ in the Calvo model. Panel (b) of [Figure A1](#) illustrates this mechanism: upstream sectors (top quartile of upstreamness) exhibit a higher fraction of adjusting firms at every horizon, and this gap widens in the months following the monetary shock as the cascade propagates. Second, cascade amplification: when an upstream firm adjusts its price, the cost change propagates through the IO matrix to downstream buyers, potentially pushing them past their own S s thresholds and triggering additional adjustments. This extensive-margin cascade is absent from Calvo models, where the fraction of adjusting firms is fixed at $1 - \theta_i$ regardless of the shock. It is also distinct from—and opposite in sign to—the intensive-margin cost-cascade channel in the linear Calvo decomposition of [Subsection 5.2](#): there, the cost channel offsets the upstreamness differential because upstream sectors have low labor shares, whereas here the extensive-margin selection effect reinforces it because threshold-crossing is more likely for the frequently-adjusting upstream sectors. The two are not in conflict; they are different margins of adjustment, and the net contribution of the menu cost extension is to raise the peak interaction modestly above the Calvo benchmark.

The two frameworks have complementary strengths. The Calvo model offers analytical tractability and clean identification of the network mechanism: the closed-form expressions in [Section 2](#) make transparent how the Leontief inverse, sectoral rigidity, and upstreamness interact to generate cross-sectional variation. With the rigidity differential disciplined externally and the demand impulse disciplined to the estimated output response, it captures 93% of the empirical peak and matches its timing. The menu cost model adds the extensive-margin selection channel and strategic complementarities, raising the peak slightly, but at the cost of one additional parameter (ψ) and the loss of closed-form expressions; its main role is as the vehicle for the asymmetry result rather than to improve the baseline fit. The impulse response is also noisier at short horizons, reflecting the state-dependent nature of the model: the exact timing of the cascade depends on the initial distribution of price gaps, which varies across Monte Carlo replications. Crucially, both models generate the same qualitative pattern: a hump-shaped interaction coefficient that peaks at medium horizons and reverts toward zero. This confirms that the network amplification mechanism documented in the main text is a robust feature of multi-sector economies with heterogeneous rigidity and IO linkages, rather than an artifact of the specific pricing friction assumed.

A-IV Data construction

This appendix provides additional details on the construction of the dataset used in the empirical analysis.

PPI series and FRED codes The Producer Price Index series are obtained from the Bureau of Labor Statistics via FRED. Each industry corresponds to a six-digit NAICS code, and the PPI series ID in FRED follows the convention PCU{NAICS}{NAICS} where NAICS is the six-digit code. The full estimation sample includes 209 manufacturing industries (NAICS 311–339) and 62 non-manufacturing industries spanning mining, utilities, trade, transportation, and services.


BEA IO table construction The IO data are from the BEA Detailed Use table, which covers industries at the six-digit NAICS level. I match BEA Detailed codes to six-digit NAICS using a cross-walk. The matching yields 271 industries with both PPI data and IO network measures used in the baseline analysis, spanning manufacturing, mining, utilities, wholesale and retail trade, transportation, and services. The upstreamness measure is computed from the full Detailed Use table to avoid understating the number of production stages. The output-allocation matrix $\mathbf{A}$ is computed by normalizing each row of the Use table by the industry’s total output, so that $a_{ik} = Z_{ik}/Y_i$ gives the share of sector i ’s output sold to sector k . The Ghosh inverse $(\mathbf{I} - \mathbf{A})^{-1}$ is then computed, and upstreamness for each industry is the row sum of this matrix. The correlation of upstreamness with eigenvector centrality is 43%, confirming that these measures capture distinct aspects of network position.

Monetary policy shock The baseline monetary policy shock is from [Bauer and Swanson \(2023\)](#). The raw data contain FOMC-meeting-level monetary policy surprises measured from high-frequency changes in federal funds futures in a 30-minute window around FOMC announcements. The “MPS_ORTH” variable orthogonalizes these raw surprises with respect to macroeconomic and financial variables to address the Fed information effect. I aggregate the FOMC-level shocks to monthly frequency by summing all surprises within each calendar month. Months without FOMC meetings receive a shock of zero. The resulting monthly series covers 1988–2023 and is matched to the panel by year-month. As an alternative measure of monetary policy shocks, I employ the dataset of [Jarociński and Karadi \(2020\)](#), which is available up to end-2024.

[Figure A2](#) displays both shock series. The Bauer-Swanson shocks exhibit the largest contractionary surprises during the 2004–2006 and 2022–2023 tightening cycles. The Jarociński-Karadi shocks are highly correlated with the Bauer-Swanson series ($\rho = 0.76$) during the overlapping period, while also capturing the large shocks of the 2022–2023 tightening cycle.⁴⁷

⁴⁷The Bauer-Swanson shock has two key advantages. First, high-frequency identification isolates the pure policy surprise at the exact moment of the announcement, avoiding the concern that monthly policy changes contain anticipated components. Second, the orthogonalization addresses the information effect directly, removing the correlation between the policy surprise and the market’s inference about the Fed’s private economic assessment.

Figure A2: Monetary policy shocks



Notes: The figure shows the monetary policy shock series: the Bauer-Swanson orthogonalized shocks and the [Jarociński and Karadi \(2020\)](#) information-purged shocks, at monthly frequency. Shaded areas indicate NBER recession dates.

A-IV.1 Summary statistics

[Table A2](#) reports summary statistics for the key variables.

Table A2: Summary statistics

Variable	N	Mean	Std. Dev.	P10	P50	P90
Monthly PPI inflation	75,073	0.002	0.018	-0.006	0.001	0.013
CPI inflation	75,073	0.002	0.003	-0.001	0.002	0.005
Bauer-Swanson orthogonalized MP surprises	71,857	0.003	0.049	-0.038	0.000	0.053
Jarociński-Karadi MP surprises	75,073	-0.004	0.047	-0.042	0.000	0.039
Upstreamness	75,073	1.222	0.352	1.009	1.103	1.566
Eigenvector centrality	75,073	0.016	0.067	0.000	0.002	0.028
Industrial production growth	75,073	0.000	0.012	-0.006	0.001	0.008
Change in Unemployment	75,073	-0.004	0.675	-0.200	0.000	0.200
M2 growth	75,073	0.005	0.006	0.000	0.004	0.009

Notes: The sample covers 271 six-digit NAICS industries over the baseline period 1998:02–2023:12 (extended to 2024:12 only for the Jarociński-Karadi robustness check). Monthly PPI inflation is the log first-difference of the industry-level producer price index. Upstreamness and eigenvector centrality are computed from BEA Detailed IO tables following [Antràs et al. \(2012\)](#). Industrial production growth is the log first-difference of the industrial production index from FRED, available at the three-digit NAICS level and assigned to all six-digit industries within each sector. Bauer-Swanson Orth. MPS is the [Bauer and Swanson \(2023\)](#) orthogonalized monetary policy surprise aggregated to monthly frequency. Jarociński-Karadi MPS is the information-purged monetary policy shock of [Jarociński and Karadi \(2020\)](#). The baseline estimation sample ends in December 2023, when the Bauer-Swanson shock series ends; the extended sample through December 2024 is used for the Jarociński-Karadi robustness check.

A-V Results

This appendix presents the full set of estimates for the results referenced in the main text.

Baseline estimates Table A3 reports the full set of baseline estimates at horizons $h = 0, 3, 6, \dots, 36$ months. Panel A reports the interaction coefficient $\hat{\beta}_h$ from the symmetric specification. Panel B reports the impulse response coefficients by upstreamness quartile. Panel C reports the asymmetric decomposition.

Real effects Table A4 reports the results about the real effects, both for the symmetric and the asymmetric specification. The specification controls for one lag of own-industry IP growth. The underlying data are based on sector-level industrial production indices from FRED for all available manufacturing sectors (NAICS 311–339). The log indices are winsorized at the 1st and 99th percentiles within each industry.

Mechanisms Table A5 reports the full set of supply chain cascade estimates at all horizons, for both the symmetric specification and the asymmetric one. Table A6 reports estimates from the triple interaction of upstreamness, intermediate input share, and the monetary policy shock. The coefficient captures whether the upstreamness effect is amplified in industries with higher intermediate input shares—a prediction of the network transmission mechanism. Table A7 reports estimates from the interaction of supplier-weighted upstreamness with the monetary policy shock. For each industry i , supplier-weighted upstreamness is defined as the input-cost-weighted average of upstream suppliers' upstreamness, computed from the BEA detailed use matrix. A significant negative coefficient indicates that industries whose suppliers are more upstream exhibit faster price adjustment, consistent with the network cascade mechanism.

Specification tests Table A8 reports the interaction coefficient $\hat{\beta}_h$ across five alternative specifications varying the lag structure, control set, and functional form. Specification A—all 271 industries with one PPI lag, zero monetary shock lags, and no macro controls—is the baseline throughout the paper.

Pre-trend Table A9 reports the estimated coefficient $\hat{\beta}_h$ of the baseline specification at negative horizons $h = -12, \dots, -1$.

Manufacturing industries Table A10 reports the set of estimates for the manufacturing sector. Both the results of the symmetric and the asymmetric specifications are reported.

Table A3: Baseline estimates

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
<i>Panel A: Symmetric specification</i>													
Upstreamness $\times$ MP shocks	-0.013 (0.010)	-0.039 (0.054)	-0.066 (0.054)	-0.059 (0.047)	-0.101*** (0.038)	-0.105** (0.041)	-0.113*** (0.039)	-0.081** (0.039)	-0.093** (0.043)	-0.097** (0.048)	-0.035 (0.037)	0.012 (0.036)	0.019 (0.039)
Observations	71,578	71,545	71,512	71,479	71,446	70,609	69,772	68,935	68,099	67,267	66,436	65,608	64,780
R^2	0.015	0.001	0.001	0.001	0.002	0.002	0.003	0.004	0.005	0.009	0.013	0.012	0.007
<i>Panel B: Estimates by upstreamness quartile</i>													
Q4 (Most Upstream)	-0.009 (0.010)	-0.027 (0.044)	-0.065 (0.053)	-0.066 (0.053)	-0.111** (0.043)	-0.120*** (0.043)	-0.127*** (0.044)	-0.117** (0.047)	-0.137*** (0.053)	-0.124* (0.064)	-0.067 (0.058)	-0.036 (0.055)	-0.027 (0.058)
Observations	16,783	16,768	16,753	16,738	16,723	16,510	16,297	16,084	15,871	15,660	15,450	15,240	15,030
R^2	0.025	0.001	0.001	0.002	0.003	0.003	0.005	0.006	0.008	0.013	0.016	0.013	0.006
Q3	-0.005 (0.005)	-0.016 (0.017)	-0.030 (0.021)	-0.030 (0.027)	-0.051** (0.025)	-0.063** (0.028)	-0.058** (0.029)	-0.051* (0.030)	-0.072* (0.040)	-0.069 (0.046)	-0.038 (0.041)	-0.021 (0.038)	-0.017 (0.040)
Observations	17,995	17,995	17,995	17,995	17,995	17,794	17,593	17,392	17,192	16,994	16,796	16,598	16,400
R^2	0.012	0.005	0.003	0.002	0.003	0.002	0.001	0.001	0.002	0.003	0.005	0.004	0.001
Q2	-0.001 (0.002)	0.004 (0.008)	0.004 (0.012)	0.002 (0.018)	-0.005 (0.019)	-0.015 (0.020)	-0.024 (0.020)	-0.024 (0.023)	-0.036 (0.027)	-0.045 (0.032)	-0.039 (0.032)	-0.032 (0.033)	-0.033 (0.033)
Observations	18,350	18,341	18,332	18,323	18,314	18,104	17,894	17,684	17,474	17,264	17,054	16,847	16,640
R^2	0.008	0.027	0.030	0.021	0.019	0.014	0.010	0.007	0.006	0.004	0.003	0.002	0.002
Q1 (Most Downstream)	0.002 (0.002)	0.003 (0.005)	0.004 (0.009)	0.007 (0.013)	-0.001 (0.015)	-0.003 (0.015)	-0.003 (0.016)	-0.007 (0.018)	-0.018 (0.021)	-0.019 (0.023)	-0.015 (0.024)	-0.014 (0.024)	-0.019 (0.027)
Observations	18,450	18,441	18,432	18,423	18,414	18,201	17,988	17,775	17,562	17,349	17,136	16,923	16,710
R^2	0.007	0.024	0.030	0.027	0.034	0.024	0.018	0.012	0.013	0.009	0.008	0.006	0.007
<i>Panel C: Asymmetric specification</i>													
Upstr. $\times$ MP shocks $\times$ exp. dummy	0.013 (0.014)	0.035 (0.028)	0.008 (0.043)	0.009 (0.059)	-0.063 (0.083)	-0.087 (0.110)	-0.158* (0.093)	-0.178* (0.091)	-0.187** (0.091)	-0.250*** (0.087)	-0.185** (0.090)	-0.142 (0.089)	-0.155* (0.093)
Upstr. $\times$ MP shocks $\times$ contr. dummy	-0.048** (0.021)	-0.140 (0.106)	-0.165 (0.103)	-0.150* (0.083)	-0.153 (0.110)	-0.128 (0.137)	-0.050 (0.152)	0.058 (0.175)	0.043 (0.181)	0.126 (0.196)	0.192 (0.189)	0.253 (0.207)	0.301 (0.224)
Observations	71,578	71,545	71,512	71,479	71,446	70,609	69,772	68,935	68,099	67,267	66,436	65,608	64,780
R^2	0.016	0.002	0.002	0.002	0.002	0.003	0.004	0.005	0.006	0.010	0.014	0.013	0.008

Notes: Each column reports local projection estimates at the indicated horizon h (months). The dependent variable is the cumulative log price change. Panel A reports the interaction coefficient from the symmetric specification in Equation 27. Panel B reports the cumulative PPI response to a one standard deviation monetary policy shock for each upstreamness quartile, estimated as four separate subsample regressions. Panel C reports estimates from the asymmetric specification in Equation 28. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A4: Estimates for industrial production

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
<i>Panel A: Symmetric specification</i>													
Upstreamness $\times$ MP shocks	-0.024* (0.014)	0.013 (0.032)	0.009 (0.024)	0.002 (0.022)	0.025 (0.020)	0.041** (0.017)	0.042** (0.017)	0.030** (0.015)	0.047*** (0.017)	0.041** (0.017)	0.046*** (0.015)	0.038* (0.022)	0.018 (0.015)
Observations	56,356	56,353	56,350	56,347	56,344	55,720	55,096	54,472	53,848	53,225	52,601	51,977	51,353
R^2	0.012	0.001	0.004	0.004	0.002	0.003	0.004	0.003	0.004	0.003	0.004	0.003	0.004
<i>Panel B: Asymmetric specification</i>													
Upstr. $\times$ MP shocks $\times$ exp. dummy	-0.040** (0.016)	0.057 (0.054)	0.019 (0.045)	0.002 (0.048)	0.005 (0.044)	-0.007 (0.032)	-0.013 (0.034)	-0.017 (0.042)	-0.007 (0.050)	-0.014 (0.049)	-0.012 (0.053)	0.006 (0.058)	-0.033 (0.051)
Upstr. $\times$ MP shocks $\times$ contr. dummy	-0.002 (0.016)	-0.049 (0.033)	-0.007 (0.027)	-0.000 (0.037)	0.053 (0.034)	0.107*** (0.036)	0.120*** (0.040)	0.097** (0.041)	0.124*** (0.048)	0.122** (0.057)	0.132*** (0.048)	0.087 (0.054)	0.100** (0.050)
Observations	56,356	56,353	56,350	56,347	56,344	55,720	55,096	54,472	53,848	53,225	52,601	51,977	51,353
R^2	0.046	0.015	0.019	0.019	0.018	0.018	0.016	0.009	0.008	0.005	0.005	0.004	0.006

Notes: Each column reports local projection estimates at the indicated horizon h (months). The dependent variable is the cumulative log industrial production growth. Panel A reports the symmetric interaction coefficient from the symmetric specification in Equation 29. Panel B reports estimates from the corresponding asymmetric specification. All specifications include industry fixed effects and one lag of IP growth. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A5: Supply chain cascade

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
<i>Panel A: Symmetric specification</i>													
Supplier price	0.176*** (0.026)	0.295*** (0.047)	0.398*** (0.078)	0.568*** (0.096)	0.684*** (0.153)	0.727*** (0.212)	0.739*** (0.204)	0.731*** (0.198)	0.665*** (0.203)	0.522** (0.253)	0.431 (0.275)	0.426 (0.296)	0.431 (0.281)
Upstreamness $\times$ MP shocks	-0.022 (0.020)	0.005 (0.048)	-0.078 (0.053)	-0.092* (0.055)	-0.154*** (0.057)	-0.184*** (0.059)	-0.158** (0.064)	-0.148** (0.067)	-0.174** (0.080)	-0.160 (0.103)	-0.037 (0.094)	0.058 (0.091)	0.091 (0.088)
Observations	29,695	29,689	29,682	29,676	29,670	29,184	28,698	28,212	27,727	27,255	26,784	26,313	25,844
R^2	0.027	0.018	0.017	0.023	0.025	0.022	0.018	0.016	0.015	0.012	0.011	0.010	0.009
<i>Panel B: Asymmetric specification</i>													
Supplier price	0.174*** (0.027)	0.296*** (0.047)	0.400*** (0.077)	0.574*** (0.094)	0.698*** (0.148)	0.753*** (0.204)	0.772*** (0.199)	0.776*** (0.193)	0.722*** (0.187)	0.597*** (0.228)	0.516** (0.246)	0.525** (0.266)	0.556** (0.261)
Upstr. $\times$ MP shocks $\times$ exp. dummy	0.033 (0.020)	0.049* (0.030)	-0.015 (0.060)	-0.102 (0.063)	-0.254*** (0.084)	-0.326*** (0.096)	-0.331*** (0.107)	-0.342*** (0.114)	-0.402*** (0.116)	-0.496*** (0.096)	-0.337*** (0.102)	-0.260*** (0.087)	-0.181*** (0.091)
Upstr. $\times$ MP shocks $\times$ contr. dummy	-0.099** (0.048)	-0.055 (0.075)	-0.165* (0.095)	-0.074 (0.133)	-0.003 (0.182)	0.030 (0.217)	0.115 (0.271)	0.175 (0.322)	0.200 (0.347)	0.395 (0.388)	0.497 (0.394)	0.650* (0.379)	0.651 (0.401)
Observations	29,695	29,689	29,682	29,676	29,670	29,184	28,698	28,212	27,727	27,255	26,784	26,313	25,844
R^2	0.029	0.019	0.018	0.024	0.028	0.028	0.025	0.025	0.027	0.030	0.032	0.034	0.033

Notes: Each column reports local projection estimates at the indicated horizon h (months). The table reports estimates from the cascade regression in Equation 31. The dependent variable is the cumulative log price change. "Supplier Price" is the input-cost-weighted average of upstream industries' PPI inflation. Panel A reports the estimates of the symmetric specification. Panel B reports the estimates of the asymmetric specification, where the interaction is split by sign of the shock. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A6: Input share amplification

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
Upstr. $\times$ int. input sh. $\times$ MP Shock	-0.009 (0.035)	-0.166 (0.197)	-0.328 (0.252)	-0.355 (0.252)	-0.466* (0.243)	-0.513** (0.222)	-0.591*** (0.221)	-0.601*** (0.217)	-0.614*** (0.200)	-0.555** (0.222)	-0.284 (0.179)	-0.122 (0.173)	-0.067 (0.190)
Observations	71,578	71,545	71,512	71,479	71,446	70,609	69,772	68,935	68,099	67,267	66,436	65,608	64,780
R^2	0.015	0.001	0.001	0.001	0.002	0.003	0.004	0.005	0.006	0.009	0.013	0.012	0.007

Notes: Each column reports local projection estimates at the indicated horizon h (months). The dependent variable is the cumulative log price change. The table reports estimates of the coefficient on the triple interaction Upstreamness $\times$ Intermediate Input Share $\times$ MP Shock. The specification controls for all lower-order interactions. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A7: Supplier-weighted upstreamness

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
Suppl. weighted upstr. $\times$ MP shocks	-0.009 (0.005)	-0.015 (0.019)	-0.031 (0.028)	-0.036 (0.035)	-0.056** (0.028)	-0.074** (0.031)	-0.079** (0.031)	-0.081** (0.034)	-0.088** (0.035)	-0.074* (0.041)	-0.039 (0.037)	-0.014 (0.035)	0.006 (0.038)
Observations	57,324	57,300	57,276	57,252	57,228	56,565	55,902	55,239	54,577	53,919	53,262	52,605	51,948
R^2	0.030	0.003	0.001	0.001	0.001	0.001	0.001	0.001	0.002	0.004	0.006	0.005	0.001

Notes: Each column reports local projection estimates at the indicated horizon h (months). The dependent variable is the cumulative log price change. The table reports estimates of the coefficient on the interaction of supplier-weighted upstreamness with the monetary policy shock, where the supplier-weighted upstreamness is the input-cost-weighted average of suppliers' upstreamness computed from the BEA detailed use matrix. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A8: Interaction coefficient across specifications

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
<i>Spec A: Baseline</i>													
Upstreamness $\times$ MP shocks	-0.013 (0.010)	-0.039 (0.054)	-0.066 (0.054)	-0.059 (0.047)	-0.101*** (0.038)	-0.105** (0.041)	-0.113*** (0.039)	-0.081** (0.039)	-0.093** (0.043)	-0.097** (0.048)	-0.035 (0.037)	0.012 (0.036)	0.019 (0.039)
Observations	71,578	71,545	71,512	71,479	71,446	70,609	69,772	68,935	68,099	67,267	66,436	65,608	64,780
R^2	0.015	0.001	0.001	0.001	0.002	0.002	0.003	0.004	0.005	0.009	0.013	0.012	0.007
<i>Spec B: 1 PPI lag, 0 MP lags, macro controls.</i>													
Upstreamness $\times$ MP shocks	-0.013 (0.010)	-0.039 (0.054)	-0.066 (0.054)	-0.059 (0.047)	-0.102*** (0.037)	-0.104** (0.041)	-0.113*** (0.039)	-0.081** (0.039)	-0.093** (0.043)	-0.096** (0.047)	-0.034 (0.037)	0.013 (0.035)	0.020 (0.038)
Observations	71,578	71,545	71,512	71,479	71,446	70,609	69,772	68,935	68,099	67,267	66,436	65,608	64,780
R^2	0.016	0.001	0.001	0.002	0.002	0.003	0.004	0.005	0.006	0.010	0.015	0.014	0.008
<i>Spec C: 6 PPI lags, 0 MP lags, no macro controls</i>													
Upstreamness $\times$ MP shocks	-0.013 (0.010)	-0.039 (0.054)	-0.064 (0.053)	-0.057 (0.046)	-0.099*** (0.037)	-0.103*** (0.040)	-0.111*** (0.038)	-0.076** (0.037)	-0.087** (0.043)	-0.087* (0.048)	-0.026 (0.038)	0.018 (0.038)	0.027 (0.037)
Observations	70,183	70,150	70,117	70,084	70,051	69,214	68,377	67,544	66,713	65,884	65,056	64,228	63,400
R^2	0.020	0.005	0.006	0.006	0.010	0.017	0.025	0.031	0.045	0.061	0.059	0.044	0.039
<i>Spec D: 1 PPI lag, 0 MP lags, $\log(\text{upstreamness})$</i>													
Upstreamness $\times$ MP shocks	-0.022 (0.017)	-0.065 (0.093)	-0.119 (0.092)	-0.107 (0.080)	-0.176*** (0.063)	-0.176*** (0.068)	-0.190*** (0.066)	-0.137** (0.067)	-0.160** (0.073)	-0.157* (0.080)	-0.051 (0.064)	0.023 (0.062)	0.033 (0.067)
Observations	71,578	71,545	71,512	71,479	71,446	70,609	69,772	68,935	68,099	67,267	66,436	65,608	64,780
R^2	0.015	0.001	0.001	0.001	0.002	0.002	0.003	0.004	0.005	0.009	0.013	0.012	0.007

Notes: Each column reports local projection estimates at the indicated horizon h (months). The dependent variable is the cumulative log price change. All specifications include industry fixed effects and sector-time fixed effects, except Spec B that features macro controls. Macro controls include CPI inflation, industrial production growth, the change in unemployment, and M2 growth. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Table A9: Interaction coefficient at negative horizons

	$h = -12$	$h = -11$	$h = -10$	$h = -9$	$h = -8$	$h = -7$	$h = -6$	$h = -5$	$h = -4$	$h = -3$	$h = -2$
Upstreamness $\times$ MP shocks	-0.004 (0.041)	0.002 (0.040)	0.001 (0.039)	-0.003 (0.037)	-0.008 (0.032)	-0.015 (0.028)	-0.020 (0.024)	-0.013 (0.020)	-0.002 (0.016)	0.002 (0.014)	-0.000 (0.010)
Observations	68,509	68,788	69,067	69,346	69,625	69,904	70,183	70,462	70,741	71,020	71,299

Notes: Each row reports the estimated coefficient on Upstreamness $\times$ MP Shock at the indicated negative horizon. The dependent variable is the cumulative log price change. Driscoll-Kraay standard errors. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. 90% confidence intervals reported. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.


Table A10: Estimates for the manufacturing sector

	$h = 0$	$h = 3$	$h = 6$	$h = 9$	$h = 12$	$h = 15$	$h = 18$	$h = 21$	$h = 24$	$h = 27$	$h = 30$	$h = 33$	$h = 36$
<i>Panel A: Symmetric specification</i>													
Upstreamness $\times$ MP shocks	-0.012 (0.012)	-0.056 (0.071)	-0.089 (0.070)	-0.078 (0.063)	-0.138*** (0.052)	-0.141** (0.057)	-0.151*** (0.052)	-0.116** (0.052)	-0.129** (0.058)	-0.132** (0.065)	-0.049 (0.050)	0.008 (0.048)	0.017 (0.052)
Observations	56,803	56,779	56,755	56,731	56,707	56,059	55,411	54,763	54,115	53,469	52,824	52,179	51,534
R^2	0.040	0.002	0.001	0.001	0.002	0.003	0.004	0.006	0.007	0.012	0.017	0.015	0.009
<i>Panel B: Asymmetric specification</i>													
Upstr. $\times$ MP shocks $\times$ exp. dummy	0.021 (0.018)	0.042 (0.038)	0.004 (0.060)	0.007 (0.080)	-0.093 (0.117)	-0.123 (0.155)	-0.215 (0.135)	-0.250* (0.130)	-0.262** (0.130)	-0.337*** (0.127)	-0.246** (0.123)	-0.192 (0.120)	-0.216* (0.122)
Upstr. $\times$ MP shocks $\times$ contr. dummy	-0.058** (0.029)	-0.188 (0.140)	-0.214 (0.141)	-0.193 (0.119)	-0.200 (0.154)	-0.165 (0.190)	-0.062 (0.209)	0.075 (0.239)	0.061 (0.245)	0.166 (0.265)	0.246 (0.256)	0.316 (0.278)	0.386 (0.301)
Observations	56,803	56,779	56,755	56,731	56,707	56,059	55,411	54,763	54,115	53,469	52,824	52,179	51,534
R^2	0.041	0.005	0.002	0.002	0.002	0.003	0.005	0.006	0.008	0.013	0.019	0.016	0.011

Notes: Each column reports local projection estimates at the indicated horizon h (months). The sample is restricted to manufacturing sectors. The dependent variable is the cumulative log price change. Panel A reports the interaction coefficient from the symmetric specification in Equation 27. Panel B reports estimates from the asymmetric specification in Equation 28. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.

Outliers Figure A3 plots the distribution of leave-one-out estimates at selected horizons. Each panel shows the distribution of the coefficient on the interaction term between upstreamness and monetary policy shocks obtained by re-estimating Equation 27 after dropping one industry from the sample. The estimates are tightly concentrated around the baseline at all horizons, confirming that no single industry drives the results.

Figure A3: Leave-one-out distribution




Notes: Each panel shows the distribution of the coefficient on the interaction term between upstreamness and monetary policy shocks from 271 leave-one-out re-estimations at the indicated horizon. The vertical dashed line marks the baseline estimate using all industries. All specifications include industry fixed effects, sector $\times$ time fixed effects, and one lag of own-industry PPI inflation.

Placebo test I randomly reassign the upstreamness measure across industries and re-estimate the baseline interaction coefficient, repeating the procedure 200 times to build a placebo distribu-

tion under the null that network position is unrelated to price responsiveness. [Figure A4](#) plots this distribution against the actual estimate at each horizon. The placebo distribution is centered near zero, whereas the baseline estimate lies far in its left tail at every horizon, indicating that the result is specific to the observed upstreamness ranking.

Figure A4: Distribution of permutation estimates



Notes: Each panel shows the distribution of the coefficient on the interaction term between upstreamness and the monetary policy shocks from 200 permutations of upstreamness across industries at the indicated horizon, along with the actual estimate from the baseline specification.

Sub-samples [Table A11](#) verifies that the main interaction coefficient is stable across sub-samples. The first column reports the full-sample baseline, while the remaining columns split the sample into the period up to the Global Financial Crisis (2010), the post-crisis period (2011 onward), and the full sample excluding the COVID period (through December 2019).

Table A11: Interaction coefficient across sub-periods

	Full sample	Up to GFC	Post- GFC	Excl. COVID-19
Upstreamness $\times$ MP shocks	-0.113*** (0.039)	-0.083* (0.048)	-0.118** (0.060)	-0.103*** (0.037)
Observations	69,772	25,592	39,622	53,795
R^2	0.003	0.011	0.004	0.004

Notes: Each column reports the estimated interaction coefficient $\hat{\beta}_{18}$ for the indicated sub-period. The dependent variable is the cumulative log price change. The pre-crisis sub-period ends in December 2010; the post-crisis sub-period begins in January 2011; the excluding-COVID sub-sample ends in December 2019. All specifications include industry fixed effects, sector-time fixed effects, and one lag of own-industry PPI inflation. Driscoll-Kraay standard errors in parentheses. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$.



PUBLICATIONS

Production Networks and the (Asymmetric) Transmission of Monetary Policy
Working Paper No. WP/2026/127