Shipping Lanes and Inflation-at-Risk:

Hub Shocks, Belief Coordination, and Optimal Monetary Policy

Prasanna Gai[1] Andrew G. Haldane[2] Olga Sudareva [3]

Abstract

We embed a hub-chain production network in a two-period New Keynesian model with dispersed beliefs and show that the structural parameters that propagate cost shocks also coordinate firms' beliefs about how persistent those shocks will be. Hub disruptions fatten the right tail of the inflation distribution because cross-sector co-movement increases the precision of firms' common signal about persistence: variance scales with the square of the network multiplier driving cost propagation, while the mean scales linearly. Panel quantile regressions across 53 economies are consistent with this asymmetry around the Russia-Ukraine conflict of 2021-22: network exposure has a near-zero linear effect on inflation but a significant, tail-concentrated quadratic effect that intensifies with supply-chain synchronisation. Expected welfare losses decompose into two components: a monetary policy amenable component, addressed by setting the interest rate above the standard prescription to internalise the future output cost of reversing immediate inflation overshoots; and a structural network welfare cost that grows quadratically with supply-chain integration and increases with the cross-sector synchronisation of supply-chain stress. A common hub shock, therefore, calls for country-specific monetary policy responses and structural resilience efforts. Calibrated to New Zealand, Australia, Germany, Korea, and Sweden, the cross-economy dispersion in network welfare costs is approximately double that in mean inflation outcomes.

[1]Gai: University of Auckland and Monetary Policy Committee, Reserve Bank of New Zealand. The views expressed herein are those of the author and do not reflect the views of the Reserve Bank of New Zealand or its Monetary Policy Committee. Email: p.gai@auckland.ac.nz.

[2]Haldane: University of Sheffield, Sheffield, United Kingdom.

[3]Sudareva: University of Auckland, Auckland, New Zealand Email: olga.sudareva@auckland.ac.nz.

Keywords: Supply chain networks, hub shocks, inflation-at-risk, belief coordination, monetary policy, network multiplier

JEL Codes: E31, E52, D85, F62

1. Introduction

The Strait of Hormuz, through which shipping lanes carry roughly a fifth of the world's traded oil and a substantial share of liquefied natural gas, is the canonical example of a hub node. Its effective closure following the Middle East conflict is the second major hub disruption to confront central banks within a single decade. In 2021–22, the disruption to global logistics was compounded by a shock to energy and commodity networks from the Russia-Ukraine war. In both episodes, the signature has been the same: costs and freight rates rising across every downstream sector simultaneously, and the distribution of inflation outcomes widening asymmetrically upward. In both, monetary policy had to be set before the persistence of the disruption was known. In the first episode, many central banks adopted a posture of treating rising inflation as "largely temporary" before belatedly embarking on a costly tightening cycle. In the current conjuncture, there is a risk that central banks may be similarly inertial: unable to determine how long the Hormuz disruption will persist, they are waiting for evidence of persistence before acting. Our paper argues that this framing, which centres monetary policy on the duration of the shock, misidentifies both the nature of the risk and the appropriate response.

The policy framework underpinning both episodes rests on a coherent foundation. The ECB's 2025 strategy review articulates the textbook prescription: look through supply shocks assessed to be temporary, and respond more forcefully only when second-round effects on wages and expectations threaten to embed the disturbance in the inflation process (Lagarde, 2025). In models where persistence is a property of the shock to be estimated and conditioned on, that prescription is defensible. But in economies organised around hub nodes, persistence is an *equilibrium outcome* of how firms form beliefs in response to the correlated signals that hub disruptions generate. A framework that treats persistence as an input – something to be learned before acting – misses the mechanism that determines it. And a policy rule that waits for persistence to be revealed will, in general, act too late: belief coordination, the process through which persistence can become self-fulfilling, will already have occurred.

Our paper links inflation risk to the topology of production networks and, in doing so, articulates this missing mechanism. This departs from network-Phillips curve models (e.g., Rubbo, 2023) where the cost-push coefficient is independent of any information structure and therefore carries no implication for the variance of inflation outcomes. The central result is an asymmetry in how hub-chain networks translate supply disruptions into inflation risk: the variance of the inflation distribution scales with the square of the network multiplier governing cost pass-through, while the mean scales only linearly. Doubling network integration doubles the mean inflation forecast but quadruples the conditional variance. Two economies receiving the same hub shock with identical mean inflation forecasts can face fundamentally different probability masses in the upper tail of the inflation distribution, depending on the depth and intensity of their supply chains. A central bank conditioning on the mean forecast alone is not merely uncertain about the future; it omits a dimension of risk whose magnitude is pinned by the architecture of production and is measurable in real time.

We take this prediction to the data. Using a panel of 53 economies from January 2015 to December 2024, panel quantile regressions around the Russia-Ukraine conflict of 2021-22 show a linear network-exposure term indistinguishable from zero at every quantile, alongside a quadratic term that is negligible at the median and rises sharply and significantly through the 75th and 90th percentiles. This is consistent with the asymmetric signature that our model predicts, and the opposite of what a purely mechanical cost-propagation account would generate.

The policy contribution of this insight turns on a distinction between what central banks can and cannot know at the time of rate-setting. Shock persistence cannot be observed before it is resolved. But central banks can observe, in real time, whether supply-chain stress indicators are moving together across sectors. This synchronisation, or the degree to which cost variation is driven by a single common source rather than sector-specific disturbances, is the structural determinant of how tightly firm beliefs coordinate around a view of persistence and, therefore, how wide the inflation distribution is likely to become. It tells a central bank how dispersed beliefs are, without knowing where those beliefs have settled.

The key implication is that pre-emptive tightening does not require knowing how long a hub disruption will last. It requires knowing how integrated the economy's supply chains are and how synchronised costs are across sectors – both observables, available from national input-output tables and standard supply-chain pressure indices. The 2021–22 episode is, in retrospect, a case study of what happens when this distinction is not drawn. Central banks in deeply integrated economies (for example, Korea or Sweden), whose calibrated network multipliers imply substantially larger network welfare costs than shallower-chain counterparts (such as Australia or New Zealand), may have applied frameworks calibrated to the mean forecast and systematically under-weighted the structural case for early tightening. The current episode offers an opportunity to explore the issue within an analytical framework that puts production networks front and centre.

Our mechanism runs through two channels that share the same structural parameters. When a disruption impairs a critical hub node, it raises unit costs in every downstream sector simultaneously, in proportion to each sector's direct and indirect dependence on hub inputs. The consumption-weighted aggregate of these pass-throughs, together with the network multiplier determined by hub intensity and chain depth, governs the mean cost-push impact on inflation. The same hub disruption also creates cross-sector cost covariance: since every sector's cost signal shares a common hub component, dispersed firms are not receiving independent information. This covariance is the structural precondition for belief coordination of the kind formalised by Morris and Shin (2002), through which pricing decisions can make persistence self-fulfilling and fatten the right tail of the inflation distribution. The precision of the common signal through which this coordination operates is amplified by the same network multiplier that governs the cost impact. The production network is simultaneously the cost-propagation mechanism and the coordination device, which is why the variance of inflation outcomes scales with the square of network integration while the mean scales only proportionally.

Our paper makes five contributions. First, we derive a *network multiplier* from a hub-chain input-output structure – the empirically dominant production architecture at the firm level (Zhao, Duley, & Gai, 2026) – and embed it in a two-period New Keynesian model with Calvo pricing and dispersed beliefs. This multiplier governs the mean cost-push, the contraction in potential output, and the rise in the natural rate of interest simultaneously, so that the monetary policy dilemma following a hub shock is a structural consequence of network topology.

Second, we show that the same multiplier governs the variance of the inflation distribution through its role in determining the precision of firms' common signal about shock persistence. The quadratic scaling of the variance relative to the linear scaling of the mean is the mechanism through which network structure generates a distributional asymmetry that mean-based frameworks omit. The framework thereby offers an analytical account of the 2021–22 episode: inflationary overshoots in deeply integrated economies were not only a function of how persistent the shocks turned out to be, but of the belief-coordination dynamics that their supply-chain topologies made structurally likely.

Third, we derive an optimal monetary policy rule that calls for greater tightening than the standard prescription. A forward-looking central bank correctly anticipates the period-2 output cost of unwinding any period-1 inflation overshoot and internalises it in the targeting rule, generating an augmentation proportional to the network multiplier and the cost-push shock. This augmented rule provides a structural foundation for the risk management tradition in monetary policy (Evans, 2019; Greenspan, 2004) – that asymmetric costs of policy errors warrant a systematic tilt away from the baseline forecast. Separately, a structural irreducible network welfare cost, which cannot be reduced by interest rate policy, grows quadratically with network integration and rises with the synchronisation of supply-chain stress indicators across sectors, providing a theoretical base for monitoring synchronisation as a first-order diagnostic variable during hub-based shock episodes.

Fourth, we embed the mechanism in a minimal open-economy setting and show that a common global hub shock generates country-specific natural rate wedges and distributional asymmetries in inflation outcomes determined by domestic supply-chain topology, producing policy divergence across countries even under an identical global disturbance. We take these cross-country implications to the data using panel quantile regressions of inflation on supply-chain synchronisation measures interacted with country-level network exposure, providing a structural test of the linear versus quadratic scaling that the theory predicts. Our results suggest that the linear term is uniformly insignificant, while the quadratic term rises sharply across the upper quantiles, with the synchronisation interaction dominating the effect at the 75th and 90th percentiles.

A final implication stems from the welfare decomposition. Since network welfare costs scale quadratically with the network multiplier, the returns to structural resilience policies that reduce hub intensity or chain depth are convex. Our model suggests that policies, including strategic inventory buffers or supply chain diversification programmes, can generate welfare gains that grow more than proportionally with the reduction in network integration they achieve.

Calibrating to input-output data for New Zealand, Australia, Germany, Korea, and Sweden, five economies spanning a wide range of supply-chain architectures, the cross-economy dispersion in network welfare costs is approximately double that in mean inflation outcomes, reflecting the quadratic versus linear relationship between welfare exposure and network integration. During the 2021–22 supply-chain episode, the optimal augmented tightening in Korea and Sweden substantially exceeds that implied by the standard rule, [by a margin identifiable from observable supply-chain topology and synchronisation data without requiring a prior assessment of shock persistence.]

Related Literature. Our paper connects three bodies of work. First, in the network macroeconomics literature, Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012) establish that network asymmetries sustain aggregate volatility; Baqaee and Farhi (2019, 2022) provide the general Leontief inverse framework for cost propagation through arbitrary networks; and Rubbo (2023) embeds network structure in the New Keynesian Phillips curve, showing that the cost-push coefficient is a consumption-weighted aggregate of sectoral pass-throughs, the analytical antecedent of our multiplier. Guerrieri, Lorenzoni, Straub, and Werning (2022) show that optimal responses to asymmetric sectoral shocks in a New Keynesian model can differ substantially from the Taylor prescription – producing demand shortages and warranting accommodative monetary policy settings. Our contribution relative to this literature is not the multiplier itself but the connection between the multiplier and the second moment of the inflation distribution, which these papers do not explore.

And the case for non-standard monetary policy in our framework arises from information-side belief coordination rather than sectoral demand spillovers and cost-side non-linearity.[4]

Second, in the strategic information literature, Morris and Shin (2002) establish that common signals are overweighted in equilibrium; our model generates a structural analogue in which the common signal is not a policy announcement but an emergent property of the input-output network, with precision disciplined by the same production parameters that govern the cost impact. This disciplines the coordination mechanism's comparative statics in a way that the sentiment-driven fluctuations of Angeletos and La'o (2013) do not. Every parameter of the belief coordination channel is tied to an observable feature of the production network.

Finally, in the inflation risk literature, Banerjee, Mehrotra, and Zampolli (2020) document that the conditional distribution of inflation is fat-tailed in ways that mean-based frameworks miss, and Adrian, Grinberg, Liang, Malik, and Yu (2022) develop the analogous growth-at-risk framework for output. We provide an analytical model of why hub shocks produce this distributional property and derive the implied policy rule from a welfare criterion rather than a reduced-form objective, giving the inflation-at-risk approach a microeconomic foundation it currently lacks.

2. Analytical Framework

2.1. Model Environment

Consider an economy composed of a *hub sector* and a sequence of *chain sectors* indexed by $j = 1, \ldots, n$. The hub sector is one whose output enters every other sector's cost structure, without exception, as an indispensable intermediate input. A critical port or shipping lane is the canonical example, being a node through which a large fraction of value-added flows. Chain sectors are

[4] See, for example, Ozhan, Sander, Wende, and Yang (2025) who develop a quantitative two-region model with occasionally binding capacity constraints.

linked sequentially. Sector j produces output x_j using three inputs: hub services, labour, and the output of its immediate upstream neighbour, sector $j-1$. Each stage of production inherits the cost pressures of every stage before it.[5] There are two periods, $s = 1, 2$.

Firms. Chain sectors are populated by a continuum of monopolistically competitive firms producing differentiated varieties under the Dixit-Stiglitz aggregator with demand elasticity $b > 1$. Monopolistic competition admits markup pricing, generates the welfare costs of price dispersion that motivate stabilisation policy, and delivers the New Keynesian Phillips curve as the equilibrium pricing condition. We assume Calvo (1983) pricing: in each period, a fraction $\gamma \in (0, 1)$ of firms cannot adjust; those able to reset solve a two-period optimisation, weighting current and expected future profits by the household stochastic discount factor.

Hub shock. At $s = 1$, suppose that a disruption leads the price of hub services to rise by $\varepsilon > 0$. Figure 1 illustrates the hub shock. Whether this shock is permanent ($\theta = 1$, reflecting a lasting geopolitical event or a permanent rerouting of global logistics) or temporary ($\theta = 0$, reflecting a resolvable bottleneck) is the central unknown. The persistence state, θ, is drawn with prior $q = \Pr(\theta = 1)$ and is common knowledge among all agents, but its realisation is not yet observed. Each firm receives noisy private and common signals about θ from its own supply-chain relationships and from the economy-wide co-movement of input costs, freight rates, and delivery times. Firms form their beliefs and set period 1 prices based on these signals.

At $s = 2$, θ is publicly revealed to all agents. If $\theta = 1$, hub costs remain elevated, period 1 prices were approximately correct, and the economy adjusts to the new cost-push steady state. If $\theta = 0$, hub costs revert to their pre-shock level, but period 1 prices were set above the period

[5] The single-chain structure is a simplification. A richer model would feature N distinct chains of heterogeneous lengths and hub intensities, each fed by the same hub. Our key findings are robust to this extension. Cost propagation along each chain follows the same geometric recursion, and the economy-wide exposure becomes a weighted average of chain-specific multipliers.

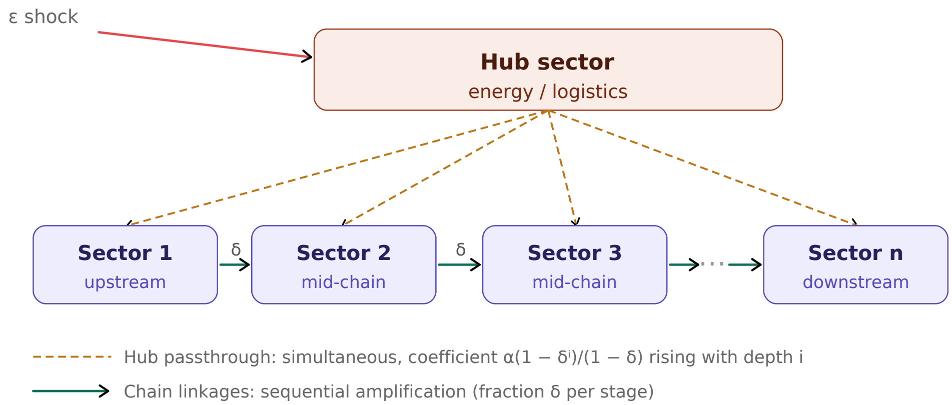

FIGURE 1. Hub-chain production network. A disruption ε at the hub enters every chain sector's unit costs simultaneously (dashed amber arrows), with a pass-through coefficient $\alpha(1-\delta^j)/(1-\delta)$ that is strictly increasing in chain depth (shown in Equation 7). Sequential linkages (teal) amplify costs stage by stage at δ per stage. A peripheral shock originating at sector k decays geometrically and does not propagate universally.

2 fundamental on the basis of beliefs that overstated persistence. This mismatch is embedded in period 2 wage contracts and input price agreements and must be unwound through a costly disinflation. The two-period structure makes precise what is meant by belief entrenchment: not merely high inflation in period 1, but a period 1 pricing error whose reversal imposes an output sacrifice that a single-period framework does not capture.

Figure 2 illustrates the timeline of the model.

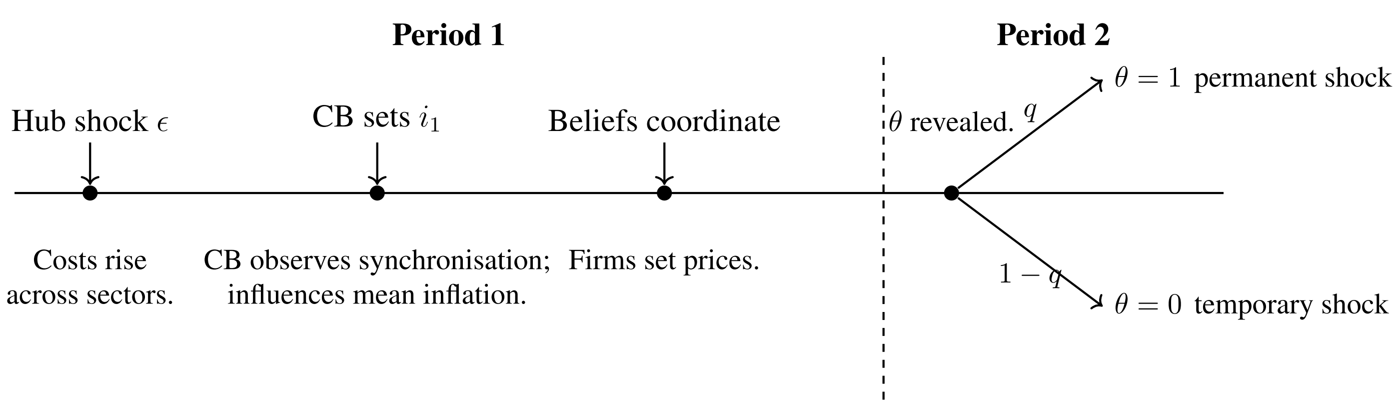

FIGURE 2. Model timing: the central bank's information problem.

Households. A representative household supplies labour and consumes the output of all sectors. It maximises expected lifetime utility over periods 1 and 2. Denoting C_s for consumption and N_s for the labour supply in period $s \in (1, 2)$, with the inverse intertemporal elasticity of substitution $\sigma > 0$, the inverse Frisch elasticity $\varphi > 0$, and the discount factor $\beta \in (0, 1)$, the household's utility is

$$U = \mathbb{E}_0 \left[\left(\frac{C_1^{1-\sigma}}{1-\sigma} - \frac{N_1^{1+\varphi}}{1+\varphi} \right) + \beta \left(\frac{C_2^{1-\sigma}}{1-\sigma} - \frac{N_2^{1+\varphi}}{1+\varphi} \right) \right], \tag{1}$$

where the expectation $\mathbb{E}_0$ is taken at the start of period 1, before signals are received and θ is realised. Optimisation yields the intra-temporal labour supply condition for each period

$$N_s^{\varphi} C_s^{\sigma} = \frac{W_s}{P_s}, \tag{2}$$

where W_s and P_s are the nominal wage and price level in period s, and the intertemporal Euler equation is

$$C_1^{-\sigma} = \beta(1 + i_1)\, \mathbb{E} \left[\frac{P_1}{P_2} C_2^{-\sigma} \right]. \tag{3}$$

Central bank. The central bank sets the nominal interest rate before the nature of the hub disruption is revealed. It observes the same aggregate supply-chain indices as firms, but cannot observe the micro-level common signal around which firm beliefs have coordinated. It chooses the period 1 nominal interest rate i_1 under this uncertainty, before θ is known.

2.2. **Production Architecture and Propagation of Hub Shocks**

The production function for the chain sector j is

$$x_j = z_j\, h^{\alpha}\, x_{j-1}^{\delta}\, l^{1-\alpha-\delta}, \tag{4}$$

where z_j is sectoral total factor productivity, $\alpha \in (0,1)$ is hub intensity, $\delta \in (0,1)$ is chain depth, $\alpha + \delta < 1$, and $x_0 \equiv 1$ so that sector 1 uses only hub services and labour. The Cobb-Douglas form implies fixed expenditure shares – α on hub services, δ on the upstream intermediate, $1-\alpha-\delta$ on labour – at any output level and any input price vector. Cost minimisation yields the unit cost

$$c_j = \frac{1}{z_j} \frac{p_h^{\alpha}}{\alpha^{\alpha}} \frac{p_{j-1}^{\delta}}{\delta^{\delta}} \frac{W^{1-\alpha-\delta}}{(1-\alpha-\delta)^{1-\alpha-\delta}}, \tag{5}$$

which, in log deviations from steady state, satisfies the recursion

$$\hat{p_j} = \alpha \hat{p_h} + \delta \hat{p_{j-1}} + (1-\alpha-\delta)\hat{w} - \hat{z_j}. \tag{6}$$

This implies that sector j's price rises in proportion to the price of hub service (with weight α, its cost share), in proportion to its upstream supplier's price (with weight δ, its intermediate cost share), in proportion to the wage (with weight $1-\alpha-\delta$), and falls with its own productivity. We normalise $\hat{p}_0 = 0$.

Cost propagation. Setting $p_h = \varepsilon$, $\hat{w} = 0$, and $z_j = 0$ for all j, Equation (6) solves forward to give

$$\hat{p_j} = \frac{\alpha(1-\delta^j)}{1-\delta}\varepsilon. \tag{7}$$

Holding the wage at its steady-state value isolates the cost-push combination of the hub shock. Three properties of Equation (7) are noteworthy.[6] First, the hub shock ε enters the price equation of every sector $j \geq 1$ without exception. Second, the pass-through coefficient $\frac{\alpha(1-\delta^j)}{1-\delta}$ is strictly increasing in j, approaching $\frac{\alpha}{1-\delta}$ as $j \to \infty$. Deeper sectors accumulate the cost burden of every

[6] Equation (7) is the Leontief inverse $(\mathbf{I} - \mathbf{A})^{-1}$ evaluated for the hub-chain topology (Baqaee & Farhi, 2019, 2022). The input-output matrix $\mathbf{A}$ places α in the hub column and δ on the sub-diagonal; inverting this lower triangular system yields the geometric series $\frac{1-\delta^j}{1-\delta}$ as the relevant element. The hub-chain motif is the empirically dominant structure in observed production networks (Zhao et al., 2026).

upstream stage. And third, an idiosyncratic disturbance at sector k enters sector $j > k$ with weight $\delta^{j-k} \to 0$ and does not propagate upstream. Hub shocks and peripheral shocks differ not merely in magnitude but in their propagation topology.

Real marginal cost. Let y denote the period 1 output gap. Using the period 1 intra-temporal condition expressed in Equation (2) to eliminate W_1 and the production function in Equation (4) to eliminate employment, real marginal cost for a firm in sector j log-deviates from the steady state by

$$\hat{mc}_j = \hat{p}_j + \omega_{mc} y_1 = \frac{\alpha(1-\delta^j)}{1-\delta}\varepsilon + \omega_{mc} y_1, \tag{8}$$

where the first term in Equation (8) is the cost-push contribution of the hub shock propagated through the chain; the second term is the standard cyclical component. And $\omega_{mc} \equiv \frac{[\sigma(1-\alpha-\delta)+\varphi]}{1-\alpha-\delta} > 1$ governs how strongly real wages respond when firms adjust employment in response to the hub shock. For $\sigma = 1, \varphi > 0$, and $\alpha + \delta < 1$, it is straightforward to have $\omega_{mc} = 1 + \frac{\varphi}{1-\alpha-\delta} > 1$.[7] A high φ (inelastic labour supply) means workers resist wage cuts, so the real wage falls slowly as employment declines; the cost-push is absorbed more into prices than into output. A low φ (elastic labour supply) means wages adjust readily, the output contraction is larger, and the inflationary impact is smaller in relative terms. Similarly, a higher σ amplifies income effects. As hub costs compress real incomes, households reduce consumption more sharply, reinforcing the output response.

Aggregating Equation (8) over sectors with CPI weights ω_j at $y_1 = 0$ yields the economy-wide cost-push term

$$\sum_j \omega_j \hat{mc}_j = \sum_j \omega_j \hat{p}_j = \Gamma(\alpha, \delta)\varepsilon, \tag{9}$$

[7] For $0 < \sigma < 1$, the condition $\omega_{mc} > 1$ additionally requires sufficiently inelastic labour supply, i.e., $\varphi > (1-\sigma)(1-\alpha-\delta)$.

where the network multiplier is defined as

$$\Gamma(\alpha, \delta) \equiv \frac{\alpha}{1-\delta} \sum_j \omega_j (1 - \delta^j). \tag{10}$$

The network multiplier, $\Gamma(\alpha, \delta)$, is the consumption-weighted aggregate pass-through of a hub shock to the economy-wide marginal cost. This scalar is a sufficient statistic for the economy's aggregate supply-side exposure to hub shocks. When $\delta = 0$, there are no chain linkages and each sector uses only hub and labour inputs, so $\Gamma = \alpha$ and the multiplier equals hub intensity with no amplification. As $\delta \to 1$, each sector inherits nearly all of its costs from upstream, so $\Gamma \to \infty$ as the geometric sums diverge. An arbitrarily deep chain with near-complete cost inheritance is highly sensitive to hub shocks. For intermediate $\delta \in (0, 1)$, $\Gamma > \alpha$, chain depth strictly amplifies the hub shock beyond its direct intensity. The leading factor $\frac{\alpha}{1-\delta}$ can be interpreted as the limiting pass-through for an infinitely long chain, while $\sum_j \omega_j (1-\delta^j)$ is the consumption-weighted average effective chain depth across sectors of the economy. Figure 3 illustrates the network multiplier.

2.3. **Aggregate Supply and Demand**

New Keynesian Phillips curve. A firm able to reset its price in period 1 chooses p_1^* to maximise the probability-weighted sum of current and expected period 2 profits. The first-order condition equates the reset price to a weighted average of current and expected future nominal marginal costs, reflecting the risk of being stuck with p_1^* in period 2. Log-linearising around the zero-inflation steady state and aggregating over resetting and non-resetting firms yields inflation as a function of average real marginal cost.

From Equation (8), average real marginal cost at $y_1 = 0$ is $\sum_j \omega_j \hat{mc}_j = \sum_j \omega_j \hat{p}_j$, which by Equation (9) and (10) equals $\Gamma(\alpha, \delta)\varepsilon$. This term enters through the input-output structure,

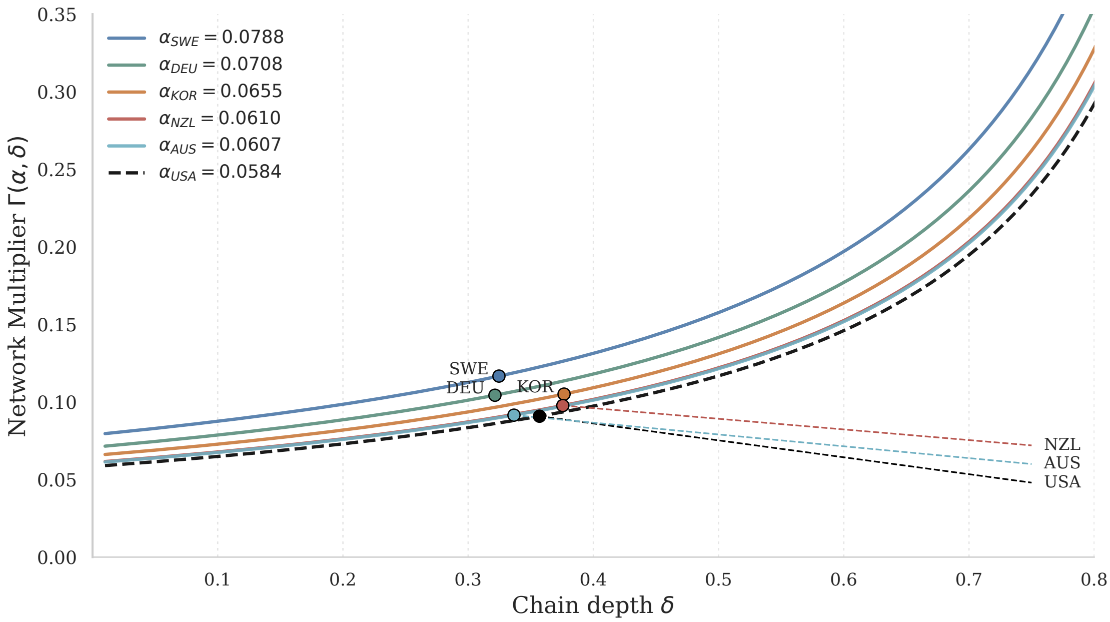

FIGURE 3. Network multiplier $\Gamma(\alpha, \delta)$ as a function of chain depth δ for six values of hub intensity α. The limiting approximation $\Gamma = \frac{\alpha}{1-\delta}$ applies for sufficiently long chains; the discrepancy from the exact finite-chain formula in Equation (10) is less than 3 per cent for the calibrated economies. Filled circles show the calibrated parameter pairs for Sweden ($\delta = 0.325$), Germany ($\delta = 0.321$), Australia ($\delta = 0.337$), United States ($\delta = 0.357$), New Zealand ($\delta = 0.376$), and Korea ($\delta = 0.377$) from Table B.1. Γ is strictly convex in δ. As supply chains become deeper (higher δ), marginal increases in network depth generate increasingly large amplification effects. Because equilibrium belief dispersion scales with Γ^2, this convexity translates into amplified higher-order systemic variance in belief coordination.

not through assumptions about firm behaviour or market power. Adding the cyclical component $\omega_{mc} y_1$, which accumulates into slope κ, and the forward-looking term delivers the New Keynesian Phillips curve in period 1:

$$\pi_1 = \kappa y_1 + \beta \mathbb{E}[\pi_2] + g\Gamma(\alpha, \delta)\varepsilon, \quad g = \frac{(1-\gamma)(1-\beta\gamma)}{\gamma}, \tag{11}$$

with slope

$$\kappa = \frac{(1-\gamma)(1-\beta\gamma)}{\gamma}\omega_{mc} = g\omega_{mc}. \tag{12}$$

The term $\beta\mathbb{E}[\pi_2]$ reflects forward-looking Calvo pricing, and $g\Gamma(\alpha, \delta)$ enters as the cost-push coefficient because it is the consumption-weighted aggregate of sectoral pass-throughs as shown in Equation (7): hub intensity (α) and chain depth (δ) determine the inflationary impact of the disruption independently of the output gap. When $\theta = 0$ in period 2 and no supply disturbance is present, the forward-looking term vanishes at the terminal condition, and the period 2's New Keynesian Phillips curve reduces to

$$\pi_2 = \kappa y_2, \tag{13}$$

from which the sacrifice ratio $\frac{1}{\kappa}$ follows as the output cost unit of period 2 disinflation.

Natural rate of output. With flexible prices, the natural rate of output for each sector is given by setting real marginal cost in Equation (8) to zero at the flexible-price equilibrium, namely:

$$y_j^* = -\frac{\alpha(1-\delta^j)}{(1-\delta)\omega_{mc}}\varepsilon. \tag{14}$$

Aggregating over sectors with CPI weights, ω_j, and using Equation (10) yields the economy-wide output:

$$y_1^* = -\frac{\Gamma(\alpha, \delta)}{\omega_{mc}}\varepsilon. \tag{15}$$

The contraction in potential output is governed by the same network multiplier as the cost-push shown in the New Keynesian Phillips curve, but is attenuated by the labour market wedge.

Period 1 IS schedule. Evaluating Equation (3), the log-linearised Euler equation, at the flexible-price equilibrium where $y_1 = 0$ and output equals the natural level, gives the hub-shocked period 1 natural real interest rate:

$$r_1^{n^*} = r^n + \frac{\Gamma(\alpha, \delta)\varepsilon}{\zeta\omega_{mc}}, \tag{16}$$

where $\zeta \equiv \frac{1}{\sigma}$. The hub-shock depresses period 1 potential output relative to period 2 and therefore increases the natural rate – a demand-side consequence that scales with the network integration and the structure of the labour market, running in parallel with the supply-side consequences in Equation (11) and (15). Substituting Equation (16) into the log-linearised Euler equation and imposing $c_1 = y_1$:

$$y_1 = -\zeta(i_1 - r^n - \mathbb{E}[\pi_2]) + \frac{\Gamma(\alpha, \delta)}{\omega_{mc}}\varepsilon. \tag{17}$$

Endogenous stagflation. Equations (11), (15) and (16) jointly imply that the hub-shock contributes a cost-push term to period 1 inflation, contracts potential output and raises the natural rate of interest, all in proportion to the same network multiplier:

$$\pi_1^{cp} = g\Gamma(\alpha, \delta)\varepsilon \tag{18}$$

$$y_1^* = -\frac{\Gamma(\alpha, \delta)}{\omega_{mc}}\varepsilon \tag{19}$$

while, from Equation (16), the natural rate rises by $\frac{\Gamma(\alpha,\delta)\varepsilon}{\zeta\omega_{mc}}$. The scale of these effects is governed by the network multiplier and deepens as hub intensity or chain depth increases. Labour market structure determines the relative magnitude. When labour supply is highly elastic or income effects are strong, the output contraction is larger relative to the price rise. This structure is a consequence of the Cobb-Douglas cost-share accounting and the household optimality conditions. Stagflation is thus endogenous.

Raising i_1 to fight the inflation in Equation (18) inflicts additional output damage on an economy whose natural rate has risen; holding i_1 low validates the inflationary pressure in Equation (18). Figure 4 illustrates the IS-NKPC equilibrium.

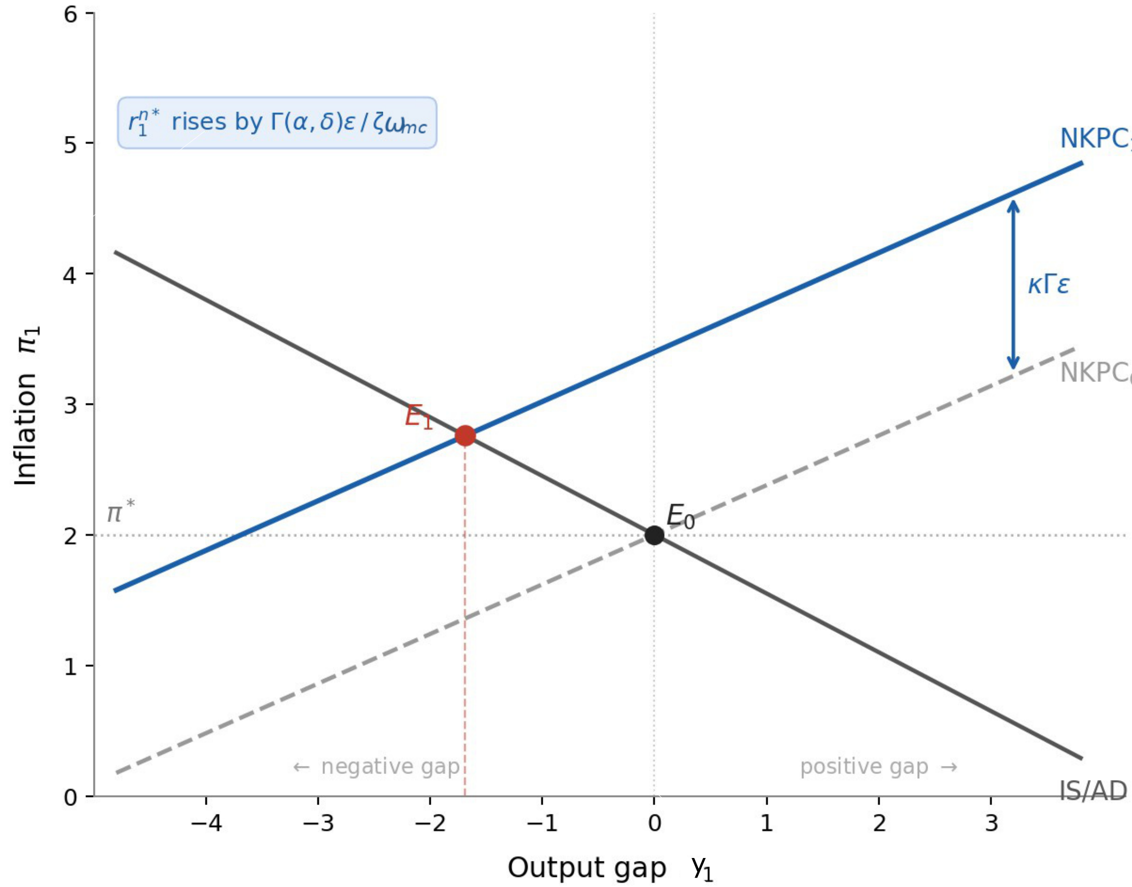

FIGURE 4. Period 1 IS-NKPC equilibrium and the stagflation dilemma. Pre-shock equilibrium lies (E_0) at the inflation target π^* and zero output gap. The hub shock ε shifts the New Keynesian Phillips curve upward by $g\Gamma(\alpha,\delta)\varepsilon$. The IS/AD schedule remains unchanged.

Cobb-Douglas as an intermediate-run approximation. A natural concern is whether the Cobb-Douglas assumption drives the result. Arguably, our assumption provides a sensible characterisation of the intermediate run, which is the horizon over which monetary policy must make its key decisions.

Intuitively, when hub prices rise, firms would like to substitute away from the now-expensive hub input toward labour or upstream intermediates. In the very short run, e.g., within a quarter, this is improbable. Production lines cannot be reengineered, supplier relationships cannot be rerouted, and energy contracts cannot be rewritten on short notice. Cost shares do not move, and so Cobb-Douglas is a reasonable approximation. In the very long run, capital can be reallocated, supply chains restructured, and alternative hubs developed, so substitution elasticities are plausibly well above unity and cost shares shift materially. But at the intermediate horizon of one to three years – the horizon relevant for monetary policy transmission – easy substitution is exhausted, and deep structural adjustment has not yet occurred. Empirical estimates of short-run input substitution elasticities in manufacturing and logistics cluster close to zero (Atalay, 2017; Boehm, Flaaen, & Pandalai-Nayar, 2019), so, if anything, Cobb-Douglas overstates the degree to which firms escape rising hub costs, making the cost pass-through in Equation (7) a lower bound on the true inflationary impact in the short run.

2.4. **Correlated Signals and the Coordination of Period 1 Beliefs**

The coordination problem. Each firm's period 1 pricing decision depends on two unknowns: the persistence of the disruption and the pricing decisions of every other firm. If all other firms absorb the shock in margins, a unilateral price increase loses customers; if all others reprice, failing to follow erodes real revenues. The optimal decision for each firm depends on what all others do. A shared belief in persistence produces the period 1 repricing that validates it; when $\theta = 0$, that repricing constitutes the pricing error whose welfare cost is derived in Section 3.2.

What makes coordination acute is that cost signals received by dispersed firms are not independent when the shock originates at the hub. From Equation (7), each sector's cost signal has a common component, θ, with sector-specific loading $\frac{\alpha^2(1-\delta^j)}{(1-\delta)}$ plus idiosyncratic noise. The cross-sector covariance between sector j and k is:

$$\mathrm{Cov}(j,k) = \frac{\alpha^2(1-\delta^j)(1-\delta^k)}{(1-\delta)^2}\varepsilon^2 q(1-q). \tag{20}$$

where $\mathrm{Var}(\theta) = q(1-q)$. This covariance is strictly positive for any $\alpha, \delta > 0$, increases monotonically in both network parameters, and equals zero in the case of peripheral shocks. Positive cross-sector covariance is the identifying signature of a hub disruption and the structural precondition for belief coordination.

Belief formation. Each firm extracts a private signal about persistence from its cost observation, $s_j = \theta + \xi_j,\ \xi_j \sim \mathcal{N}(0, \sigma_s^2)$ where σ_s^2 captures the effective rescaling by the loading factor $\frac{\alpha^2(1-\delta^j)^2}{(1-\delta)^2}$. Firms also observe a common signal, $v = \theta + \eta,\ \eta \sim \mathcal{N}(0, \sigma_v^2)$, extracted from the economy-wide co-movement of supply-chain indicators. Appendix A derives $\frac{1}{\sigma_v^2} = K\Gamma^2(\alpha, \delta)\varepsilon^2$, where $K \equiv \frac{N_{eff}}{\sigma_s^2}$ measures how much information the cross-sectional average of firm cost signals contains about the true persistence state θ. $N_{eff} = \frac{1}{\sum_j \omega_j^2}$ is defined as the reciprocal of the Herfindahl-Hirschman concentration measure of sectoral CPI weights.

Under the linear-normal approximation (Morris & Shin, 2002), the average firm belief about persistence is

$$\mu_j \approx (1-\rho)s_j + \rho v, \quad \rho = \frac{1/\sigma_v^2}{1/\sigma_v^2 + 1/\sigma_s^2}. \tag{21}$$

Proposition 1. *Under the linear-normal approximation,* $\mathrm{Var}(\mu_j) = (1-\rho)^2\sigma_s^2$ *is strictly decreasing in* ρ.

Proof. Conditional on (θ, v), the s_j are independent with variance σ_s^2, giving $\mathrm{Var}(\mu_j \mid \theta, v) = (1-\rho)^2\sigma_s^2$. Differentiating in ρ gives $-2(1-\rho)\sigma_s^2 < 0$. □

As $\rho \to 1$, all firms form nearly identical posteriors and period 1 beliefs coordinate tightly around v.

The relative precision, ρ, measures how much weight each firm places on the common signal relative to its own private experience. When ρ is high, firms are effectively reading from the same page, and the average posterior μ_j becomes a tight reflection of whenever the common signal v has landed. The network multiplier, $\Gamma(\alpha, \delta)$, governs how strongly the network amplifies the hub shock into cross-sector cost covariance: deeper chains and more hub-intensive production mean more of each sector's cost experience is driven by the same upstream source, making the signals across sectors more correlated and the common signal more informative.

The term $K \equiv \frac{N_{eff}}{\sigma_s^2}$ captures the signal quality of the aggregate supply-chain index, since N_{eff} quantifies how much idiosyncratic noise is washed out when the central bank or firms average across sectoral indicators. An economy with many roughly equal-weighted sectors has a high N_{eff} and a precise aggregate signal; one dominated by a handful of large CPI categories has a low N_{eff} and a noisy aggregate signal.

In sum, network architecture, $\Gamma(\alpha, \delta)$, determines how correlated sectoral cost signals are; that correlation, scaled by the signal quality K, determines how precisely the common signal v is identified; and the precision of v relative to private experience determines ρ, the degree to which firms abandon their private information and coordinate on a shared view of persistence.

Corollary 1. *The relative precision satisfies*

$$\rho(\alpha, \delta, \varepsilon) = \frac{K\Gamma^2(\alpha, \delta)\varepsilon^2}{K\Gamma^2(\alpha, \delta)\varepsilon^2 + 1/\sigma_s^2}, \tag{22}$$

which is strictly increasing in α, δ, *and* ε.

Proof. Substituting $1/\sigma_v^2 = K\Gamma^2(\alpha, \delta)\varepsilon^2$ into the definition of ρ gives Equation (22). Monotonicity in ε follows from the quotient rule applied to a strictly increasing numerator. Monotonicity in α from $\frac{\partial \Gamma}{\partial \alpha} > 0$. And monotonicity in δ follows from $\frac{\partial \Gamma}{\partial \delta} > 0$ since $\frac{1-\delta^j}{1-\delta} = 1 + \delta + \cdots + \delta^{j-1}$ is strictly increasing in δ for $j \geq 2$. $\square$

The parameters α and δ that determine the cost impact through $\Gamma(\alpha, \delta)$ in Equation (18) and (19) are precisely those that determine belief coordination through $\rho(\alpha, \delta, \varepsilon)$ in Equation (22). The hub-chain network is both the cost-propagation mechanism and the coordination device. A more hub-intensive, deeper-chain economy is simultaneously more exposed to cost pass-through and more susceptible to belief entrenchment.

Period 1 equilibrium prices. Each firm sets its reset prices as a weighted average of its own expected costs and the expected aggregate price level. Under the standard properties of Calvo pricing first-order condition and aggregating across firms, the period 1 aggregate log price level is

$$p_1 = \mu\Gamma(\alpha, \delta)\varepsilon, \quad \mu = (1-\rho)\theta + \rho v. \tag{23}$$

where μ is the cross-sectional average belief about persistence. Since both μ and Γ depend on α and δ, the effect of deeper integration on period 1 prices decomposes as

$$\frac{\partial p_1}{\partial \alpha} = \underbrace{(v-\theta)\rho_\alpha \Gamma\varepsilon}_{\text{belief coordination channel}} + \underbrace{\mu\frac{\partial \Gamma}{\partial \alpha}\varepsilon}_{\text{cost amplification channel}} . \tag{24}$$

Lemma 1. *Suppose the cross-sectional average belief satisfies $\mu > 0$, then: (i). the cost-amplification channel, $\mu\frac{\partial\Gamma}{\partial\alpha}\varepsilon$, is strictly positive; (ii). the belief coordination channel, $(v-\theta)\rho_\alpha\Gamma\varepsilon$, has sign $(v-\theta)$; (iii). when $v \geq \theta$, $\frac{\partial p_1}{\partial\alpha} > 0$; (iv). when $v < \theta$, $\frac{\partial p_1}{\partial\alpha} > 0$ if and only if $\mu\frac{\partial\Gamma}{\partial\alpha} > (\theta - v)\rho_\alpha\Gamma$. The analogous statement holds for δ.*

Proof. (i). Under the assumption $\mu > 0$, with $\frac{\partial\Gamma}{\partial\alpha} > 0$ from Equation (10); (ii). Since $\rho_\alpha > 0$ and $\Gamma\varepsilon > 0$, the sign of the belief coordination channel equals the sign of $(v-\theta)$; (iii). When $v \geq \theta$, the belief coordination channel is non-negative by (ii) and the cost amplification channel is positive by (i). So their sum is strictly positive; (iv). When $v < \theta$, (ii) gives a strictly negative belief coordination channel and (i) gives a strictly positive cost amplification channel. So, the sum is positive if and only if the latter dominates in the absolute value. □

When $v \geq \theta$ – the common signal overstates persistence – both channels amplify the period 1 pricing error and the period 2 disinflation cost derived in Section 3. A hub-intensive economy does not uniformly generate higher inflation; it generates fatter tails in both directions, with the outcome determined by where v falls relative to θ.

3. Fat Tails, Welfare, and Monetary Policy

We now establish the paper's central results. Section 3.1 derives the distribution of inflation implied by the hub-shock mechanism and shows that the variance scales with the square of the network multiplier while the mean scales only linearly. Section 3.2 demonstrates that this variance term enters welfare through the expected cost of belief-driven pricing errors and derives the optimal policy rule that corrects for it. Section 3.3 shows that belief coordination also raises the effective sacrifice ratio for period 2 disinflation, reinforcing the case for pre-emptive tightening. Together, the three results establish that optimal monetary policy requires tightening beyond the standard presumption by an amount proportional to the network multiplier. Furthermore, expected welfare includes a network welfare cost, $\mathcal{F}(S)$, that monetary policy cannot eliminate.

3.1. Network Stress and the Inflation Distribution

Network stress. The central bank observes the synchronisation of supply-chain indicators across sectors, but not the realisation of the common signal v around which firm beliefs have converged.

Define period 1 network stress as the signal-to-noise ratio of the common hub component in sectoral cost signals:

$$S \equiv \frac{K\Gamma^2(\alpha,\delta)\varepsilon^2}{K\Gamma^2(\alpha,\delta)\varepsilon^2 + 1/\sigma_s^2}. \tag{25}$$

When S is high, most of the cross-sector variation in costs reflects a single common source rather than sector-specific disturbances. A central bank observing that freight rates, energy prices, and delivery times are all moving together is observing high S: the synchronisation itself is the signal, not the level of any individual indicator. From Proposition 1 and Corollary 1, $S = \rho(\alpha, \delta, \varepsilon)$ exactly, so S rises monotonically in hub intensity, chain depth, and shock magnitude. Aggregate supply-chain indices, such as the GSCPI, are empirical proxies for S. Observing S reveals the width of the posterior distribution over persistence – how tightly firm beliefs are bunched – but nothing about where the centre lies. This gap generates the central bank's residual uncertainty about μ.

The inflation distribution. Period 1 inflation is

$$\pi_1 = \mu\Gamma(\alpha,\delta)\varepsilon + u, \quad u \sim \mathcal{N}(0, \sigma_u^2), \tag{26}$$

where u captures demand fluctuations orthogonal to the hub-shock. Decomposing $\mu = \mu^* + \upsilon$, where μ^* is the central bank's estimate of average persistence beliefs and υ is its forecast error, the distribution of inflation conditional on S has mean $\overline{\pi} = \mu^*\Gamma\varepsilon$ and variance

$$\sum{}^2(S) = \Gamma^2(\alpha,\delta)\varepsilon^2\tau^2(S) + \sigma_u^2, \tag{27}$$

where $\tau^2(S) \equiv \mathrm{Var}(\upsilon \mid S)$. The belief coordination uncertainty scales with $\Gamma^2(\alpha, \delta)$ rather than $\Gamma(\alpha, \delta)$ because the cost gap $\Gamma(\alpha, \delta)\varepsilon$ multiplies the uncertain belief component υ, and variances scale with the square of multipliers. Applying Gaussian conjugate updating and adopting the linear approximation $\tau^2(S) \approx \tau_0^2(S)$ gives the closed-form expression used throughout.[8]

$$\sum{}^2(S) = \Gamma^2(\alpha, \delta)\varepsilon^2\tau_0^2(S) + \sigma_u^2. \tag{28}$$

The 90th-percentile inflation-at-risk is

$$\pi^{ATR} = \overline{\pi} + 1.28\sum(S). \tag{29}$$

Proposition 2. *The mean of the inflation distribution scales linearly with the network multiplier, $\overline{\pi} \propto \Gamma$, while the conditional variance scales quadratically, $\sum{}^2(S) \propto \Gamma^2$. Doubling network integration doubles the mean inflation forecast but quadruples the conditional variance. Two economies with identical $\overline{\pi}$ but different network architectures face materially different probability masses in the upper tail.*

Proof. The mean $\overline{\pi} = \mu^*\Gamma\varepsilon$ is linear in Γ by inspection of Equation (26). The variance in Equation (27) contains $\Gamma^2(\alpha, \delta)$ as the coefficient of the belief-coordination term: $\Gamma\varepsilon$ multiplies υ in Equation (26) so that $\mathrm{Var}(\Gamma\varepsilon\,\upsilon \mid \mathrm{S}) = \Gamma^2(\alpha, \delta)\varepsilon^2\tau^2(S)$. □

Figure 5 illustrates the result. At low network stress, the conditional distribution is narrow and roughly symmetric; at high network stress, the right tail fattens disproportionately while the mean shifts only modestly. A central bank conditioning only on $\overline{\pi}$ is blind to this asymmetry. The welfare cost of that blindness is the subject of Section 3.2.

[8] The linear approximation $\tau^2(S) \approx \tau_0^2(S)$ is a leading-order Taylor expansion that delivers the closed-form variance in Equation (27). All qualitative results in what follows depend on $\frac{\partial \Sigma^2}{\partial S} > 0$, which holds for strictly increasing $\tau^2(S)$.

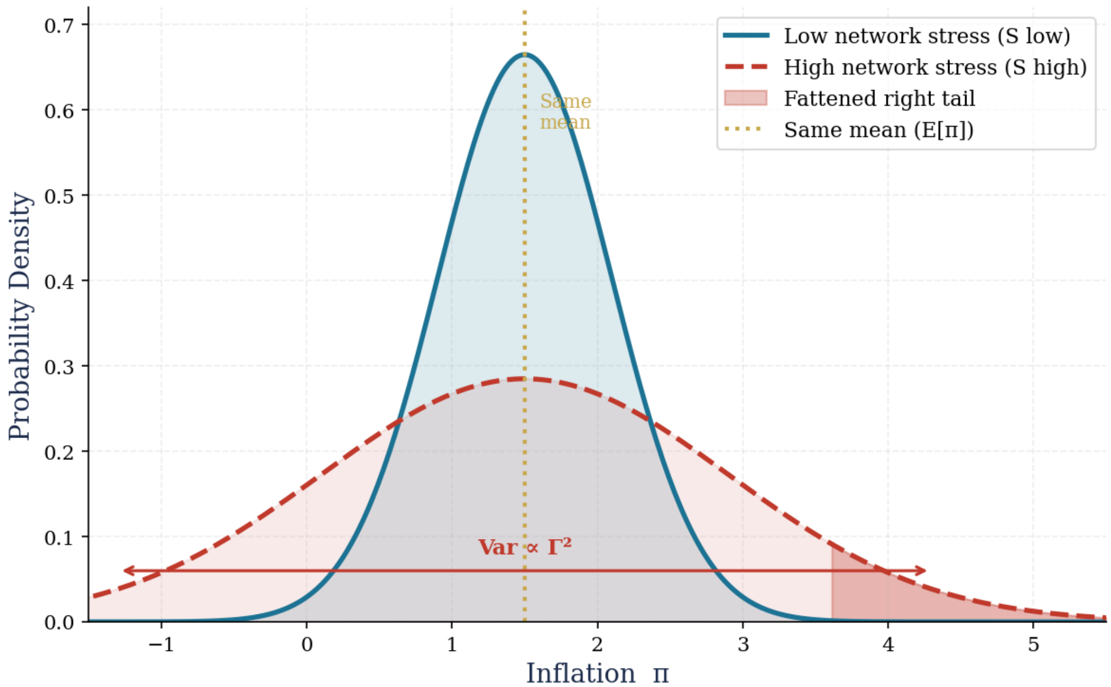

FIGURE 5. The inflation distribution and inflation-at-risk. Under low network stress (in the blue line), the conditional distribution of period 1 inflation is narrow and roughly symmetric around the mean of 1.5 per cent. Under high network stress (in the red dashed line), the distribution fattens its tail disproportionately. Two economies with identical mean inflation forecasts but different network architectures face materially different probability masses in the upper tail.

3.2. **Welfare Foundations and the Missing Variance Term**

We now set up the two-period welfare criterion and show that it contains a term proportional to $\Gamma^2(\alpha, \delta)\varepsilon^2\tau^2(S)$ that the standard loss function omits. This omitted term is structural in the production network, and its magnitude is the basis for the policy gap, Δi, and the network welfare cost, $\mathcal{F}(S)$, derived in Section 3.3.

Period 1 welfare loss. Following Woodford (2003) and Benigno and Woodford (2005), a second-order approximation to lifetime utility in Equation (1) around the efficient allocation gives, under a production subsidy that eliminates the steady-state markup:

$$\mathcal{W}_1 = \frac{1}{2}\mathbb{E}_0[(\pi_1 - \pi^*)^2 + \lambda y_1^2], \quad \lambda \equiv \frac{\kappa}{b}. \tag{30}$$

Both κ and $\lambda \equiv \frac{\kappa}{b}$ are pinned by structural parameters. This is the standard criterion. The central bank's blind spot arises not from misspecifying the welfare in Equation (30) but from treating it as the complete welfare criterion when it is only the first of two terms.

Period 2 disinflation cost. When $\theta = 0$, hub costs revert, and the period 2 New Keynesian Phillips curve is $\pi_2 = \kappa y_2$, as shown in Equation (13). Period 1 prices were set at $p_1 = \mu\Gamma(\alpha, \delta)\varepsilon > 0$ on the basis of period 1 beliefs from Equation (23) that overstated persistence. Sticky wages and indexed input contracts carry this error into period 2, so unwinding the price-level deviation requires an output sacrifice

$$y_2^D = -\frac{\pi_1}{\kappa}, \tag{31}$$

and applying the second-order welfare approximation to period 2 utility, including both inflation deviation and output sacrifice, gives

$$\mathcal{W}_2(\theta = 0) = \frac{\kappa^2 + \lambda}{2\kappa^2}\pi_1^2. \tag{32}$$

This cost is quadratic in the network multiplier: deeper, more hub-intensive networks generate larger pricing errors and bear a disproportionately heavier welfare cost.

The missing term. Combining Equation (30) and (32) over both periods, with β discounting period 2 and $(1 - q)$ the probability that disinflation is required:

$$\mathbb{E}[\mathcal{W}] = \mathcal{W}_1 + \frac{\beta(1 - q)(\kappa^2 + \lambda))}{2\kappa^2}\mathbb{E}[\pi_1^2]. \tag{33}$$

Defining $\psi_c = \frac{\beta(1-q)(\kappa^2+\lambda)}{\kappa^2}$ and using the decomposition $\mathbb{E}[\pi_1^2] = \overline{\pi}^2 + \Sigma^2(S)$ with $\Sigma^2(S) = \Gamma^2(\alpha,\delta)\varepsilon^2\tau^2(S) + \sigma_u^2$ from Section 3.1, the period 2 contribution to expected welfare is

$$\mathbb{E}[\mathcal{W}_2] = \frac{\psi_c}{2}\mathbb{E}[\pi_1^2] = \frac{\psi_c}{2}[\overline{\pi}^2 + \Gamma^2(\alpha,\delta)\varepsilon^2\tau^2(S) + \sigma_u^2]. \tag{34}$$

The term $\Gamma^2(\alpha,\delta)\varepsilon^2\tau^2(S)$ is entirely absent from the standard loss function $\mathcal{W}_1$. It captures the expected cost of belief-driven pricing errors that must be unwound when firms overestimate the persistence of the shock. All parameters of the augmented loss are pinned by the structural parameters of the model.

3.3. **Optimal Monetary Policy**

Section 3.2 establishes that the standard welfare criterion $\mathcal{W}_1$ omits an expected cost $\frac{\psi_c}{2}\mathbb{E}[\pi_1^2]$, namely the period-2 disinflation burden generated when period-1 beliefs overshoot and firms price "as if" the hub disruption will persist. In this section, we derive the implications for optimal monetary policy. The omitted cost modifies the central bank's effective inflation-aversion, tightening the optimal policy rule above the standard prescription. We then decompose total welfare into a component that monetary policy can reduce, and an irreducible component, a network welfare cost $\mathcal{F}(S)$, that is an object of structural policy and whose magnitude is determined by the architecture of production.

The standard rule. A central bank optimising $\mathcal{W}_1$ alone, subject to NKPC (Equation 11) and IS (Equation 17) sets its optimal targeting rule by equating the marginal welfare cost of inflation with the marginal output cost of reducing it:

$$\kappa(\pi_1 - \pi^*) + \lambda y_1 = 0. \tag{35}$$

Solving jointly with the NKPC and recovering the interest rate from the IS schedule and the hub-shocked natural rate in Equation (16) yields the standard rule:

$$i_Q^* = r^n + \overline{\pi} + \underbrace{\frac{\kappa g \Gamma(\alpha, \delta)\varepsilon}{\zeta(\kappa^2 + \lambda)}}_{\text{cost-push response}} + \underbrace{\frac{\Gamma(\alpha, \delta)\varepsilon}{\zeta \omega_{mc}}}_{\text{natural rate adjustment}} . \tag{36}$$

The third term is the standard NK response to the cost-push shock $g\Gamma(\alpha, \delta)\varepsilon$ in Equation (11). The fourth term raises the policy rate to match the hub-shocked natural rate from Equation (16). The rule correctly recognises the stagflationary configuration of Section 2.3 but is derived from $\mathcal{W}_1$ alone and is, therefore, blind to the period-2 inflation cost.

Augmented welfare and optimal monetary policy. The full expected welfare from Equations (33)-(34) is

$$\mathbb{E}[\mathcal{W}] = \mathcal{W}_1 + \frac{\psi_c}{2}\mathbb{E}[\pi_1^2], \tag{37}$$

where the expected squared period-1 inflation decomposes as

$$\mathbb{E}[\pi_1^2] = \overline{\pi}^2 + \sum{}^2(S), \quad \sum{}^2(S) = \Gamma^2(\alpha, \delta)\varepsilon^2\tau_0^2(S) + \sigma_u^2. \tag{38}$$

This decomposition means that the central bank can reduce $\overline{\pi} = \mu^*\Gamma(\alpha, \delta)\varepsilon$ through interest rate policy rate via the IS-NKPC transmission, but it cannot reduce $\sum{}^2(S)$, which is determined entirely by the production network architecture and the degree of supply-chain synchronisation. The two components, therefore, call for distinct treatments:

- The reducible component $\overline{\pi}^2$ creates an additional incentive to tighten. The central bank can lower it by raising i_1, and the optimal rule accounts for this through a modified targeting rule.

- The network welfare cost $\sum{}^{2}(S)$ is a welfare burden the central bank faces regardless of its rate decision. It cannot be eliminated through interest rate policy, but it must be recognised: a central bank that ignores it systematically understates the welfare cost of hub disruptions and prices the case for pre-emptive tightening too cheaply.

Differentiating $\mathbb{E}[\mathcal{W}]$ with respect to i_1 and using the structural derivatives $\frac{\partial \pi_1}{\partial i_1} = -\kappa\zeta$ and $\frac{\partial y_1}{\partial i_1} = -\zeta$ from the IS-NKPC system, the first-order condition is

$$-\kappa\zeta[(\pi_1 - \pi^*) + \psi_c \pi_1] - \lambda\zeta y_1 = 0. \tag{39}$$

Dividing through by ζ yields the targeting rule:

$$\kappa[(1 + \psi_c)\pi_1 - \pi^*] + \lambda y_1 = 0. \tag{40}$$

Compared to the standard targeting rule, the augmented version places weight $(1 + \psi_c)$ on π_1 rather than 1: the central bank is more than inflation-averse because the period-1 inflation above π^* will need to be unwound in period 2 at an output cost, and the central bank internalises that future sacrifice. The optimal rule is therefore

$$i_N^* = r^n + \overline{\pi}^N + \underbrace{\frac{\kappa(1 + \psi_c)g\Gamma(\alpha, \delta)\varepsilon}{\zeta(\kappa^2(1 + \psi_c) + \lambda)}}_{\text{augmented cost-push response}} + \underbrace{\frac{\Gamma(\alpha, \delta)\varepsilon}{\zeta\omega_{mc}}}_{\text{natural rate adjustment}}, \tag{41}$$

where $\overline{\pi}^N = \frac{\lambda g\Gamma(\alpha,\delta)\varepsilon}{\kappa^2(1+\psi_c)+\lambda}$.

The structure mirrors the standard rule in Equation (36) with two changes: The cost-push response coefficient increases from $\frac{\kappa}{\kappa^2+\lambda}$ to $\frac{\kappa(1+\psi_c)}{\kappa^2(1+\psi_c)+\lambda}$: the central bank responds more aggressively to the cost-push shock because inflation is more costly when future disinflation burdens are internalised. The Fisher equation component uses $\overline{\pi}^N$ rather than $\overline{\pi}$, reflecting lower equilibrium inflation under the tighter targeting rule. The natural rate adjustment term is unchanged.

The policy gap $\Delta i = i_N^* - i_Q^*$ is obtained by subtracting Equation (36) from Equation (41). The natural rate adjustment terms cancel identically. We have

$$\Delta i = \frac{\psi_c \lambda \kappa g \Gamma(\alpha, \delta)\varepsilon(1 - \kappa\zeta)}{\zeta(\kappa^2 + \lambda)[\kappa^2(1 + \psi_c) + \lambda]}. \tag{42}$$

For first-order accuracy in ψ_c (appropriate when the period-2 cost is not too large relative to the period-1 welfare), this simplifies to

$$\Delta i \approx \frac{\psi_c \lambda \kappa g \Gamma(\alpha, \delta)\varepsilon(1 - \kappa\zeta)}{\zeta(\kappa^2 + \lambda)}. \tag{43}$$

The policy gap is positive when $\kappa\zeta < 1$. The NKPC slope $\kappa = g\omega_{mc}$ is typically in the range $0.05 - 0.30$ in estimated NK models, and the intertemporal elasticity $\zeta = \frac{1}{\sigma}$ is typically $0.5 - 1.0$ for $\sigma \in [1, 2]$. We, therefore, treat $\kappa\zeta < 1$ as the empirically relevant case. Deeper or more hub-intensive networks, and larger shocks, call for greater tightening above the standard rule.

Network welfare cost. Even under the optimal rule i_N^*, total expected welfare contains a term that monetary policy cannot eliminate:

$$\mathcal{F}(S) \equiv \frac{\psi_c}{2} \sum{}^2(S) = \frac{\psi_c}{2}[\Gamma^2(\alpha, \delta)\varepsilon^2 \tau_0^2(S) + \sigma_u^2]. \tag{44}$$

$\mathcal{F}(S)$ is the expected disinflation burden attributable to the variance of the inflation distribution – the part driven by the central bank's residual uncertainty about where firm beliefs have settled, given observed network stress S. It is strictly positive at all $S \geq 0$ (since $\sigma^2 > 0$), strictly increasing in S (since $\tau_0^2(S)$ is increasing with $\tau_0^2(0) = 0$ from footnote 8), and scales quadratically with $\Gamma(\alpha, \delta)$. A central bank optimising $\mathcal{W}_1$ alone incurs the cost $\mathcal{F}(S)$; and the cost grows quadratically as network integration deepens.

Proposition 3. *Under the condition $\kappa\zeta < 1$, the policy gap satisfies (i). strict positivity, $\Delta i > 0$, for all $\psi_c > 0$, $\Gamma(\alpha, \delta) > 0$, $\varepsilon > 0$; (ii). monotonicity in network parameters, Δi is strictly increasing in $\Gamma(\alpha, \delta)$ and ε; (iii). monotonicity in disinflation-cost weight and intertemporal elasticity, Δi is strictly increasing in ψ_c and is strictly decreasing in ζ; (iv). The network welfare cost $\mathcal{F}(S)$ is strictly positive for all $S \geq 0$, strictly increasing in S; (v). and scales quadratically with $\Gamma(\alpha, \delta)$. It is invariant to the choice of i_1 and cannot be reduced by interest rate policy.*

Proof. (i). All factors in Equation (42) are positive under $\kappa\zeta < 1$. (ii). $\frac{\partial \Delta i}{\partial \Gamma} = \frac{\Delta i}{\Gamma} > 0$; $\frac{\partial \Delta i}{\partial \varepsilon} = \frac{\Delta i}{\varepsilon} > 0$; monotonicity in α and δ follows from $\frac{\partial \Gamma(\alpha,\delta)}{\partial \alpha} > 0$ and $\frac{\partial \Gamma(\alpha,\delta)}{\partial \delta} > 0$. (iii). Direct differentiation of Equation (42). (iv). $\frac{\partial \Delta i}{\partial \zeta} = -\frac{\Delta i}{\zeta} < 0$. (v). Positivity from $\sigma_u^2 > 0$; monotonicity in S from $\frac{\partial \tau_0^2(S)}{\partial S} > 0$; quadratic scaling from $\Gamma^2(\alpha, \delta)\varepsilon^2\tau_0^2(S)$ term; invariance to i_1, because $\sum{}^2(S)$ has no i_1-dependence. $\square$

Deeper or more hub-intensive networks, and larger shocks, call for greater tightening above the standard rule. The standard rule i_Q^* implicitly sets $\psi_c = 0$, ignoring the period-2 disinflation cost entirely. The optimal rule i_N^* raises the interest rate by $\Delta i > 0$ through a more inflation-averse targeting rule. Two episodes with identical $\overline{\pi}$ but different $\sum{}^2(S)$ receive identical treatment under i_Q^*; under i_N^*, both receive the same rate adjustment Δi, but the episode with higher S faces a larger irreducible network welfare cost $\mathcal{F}(S)$ that optimal policy cannot eliminate. A central bank conditioning on the mean forecast alone, therefore, systematically understates the welfare cost of hub disruptions by the amount $\mathcal{F}(S)$, which grows quadratically with network integration and is increasing in supply-chain synchronisation S.

Risk management perspective. The policy gap Δi is proportional to $g\Gamma(\alpha, \delta)\varepsilon$ (the cost-push shock that shifts the NKPC) and to ψ_c (the structural weight on future disinflation costs). It is not strictly proportional to $\sum{}^{2}(S)$. What $\sum{}^{2}(S)$ determines is the network welfare cost $\mathcal{F}(S)$ that motivates setting $\psi_c > 0$ in the first place: the larger the irreducible risk exposure, the stronger the welfare case for the central bank to internalise disinflation costs. This is structurally consistent with the risk-management tradition articulated by Greenspan (2004) and formalised by Evans (2019): when the distribution of outcomes is fat-tailed, and the tail is structural rather than random, optimal policy should account for the asymmetric costs, even when the precise outcome is uncertain.

Our model provides the risk management approach with an explicit welfare foundation. The network welfare cost $\mathcal{F}(S)$ grows quadratically with network integration, while the mean cost grows only linearly; and the optimal policy response – embodied in $\Delta i > 0$ – ensures that the reducible portion is internalised. Three practical conclusions follow. First, when a hub disruption occurs, the appropriate benchmark for rate-setting is the hub-shocked natural rate, not the pre-shock neutral. Second, the key monitoring object is $\mathcal{F}(S)$ and its determinants – the synchronisation of stress indicators across sectors and the depth of supply chains —- not the level of aggregate supply-chain indices alone. Third, the time to respond is before the tail has fully fattened, because $\mathcal{F}(S)$ is irreducible once belief coordination has set in: the cost of waiting is borne permanently.

Resilience policies and convex returns. Since $\mathcal{F}(S)$ has no dependence on monetary policy, the burden of a reduction in network welfare cost falls in the domain of fiscal and structural policy.

Three instruments follow directly from the model parameters. First, investment in alternative hub capacity, including domestic energy production, reduces hub intensity, α, lowering the cost-push impact and the precision of belief coordination. Second, supply-chain diversification programmes that shorten or duplicate critical chain stages compress the network multiplier and therefore welfare costs, $\mathcal{F}(S)$, quadratically. And third, strategic inventory buffers dilute the magnitude of the hub shock, ε, experienced by downstream sectors, lowering S and, alternatively, the right tail of the inflation distribution. Since $\mathcal{F}(S)$ is proportional to $\Gamma^2(\alpha, \delta)$, a reduction in Γ generates a welfare gain that is proportionally larger the higher the starting level of integration.

4. GLOBAL HUB SHOCKS AND OPEN-ECONOMY IMPLICATIONS

We now embed our mechanism in a minimal open-economy setting. Two results emerge. First, a common global hub shock generates cross-country interest rate divergence through three channels: natural rate wedges, mean inflation differentials, and disinflation-cost premia that all scale *linearly* in the domestic network multiplier Γ_m. Second, the network welfare costs diverge *quadratically* across countries: a country with twice the network multiplier faces four times the irreducible welfare burden, regardless of its monetary policy response. The variance of the inflation distribution, and the welfare cost associated with it, is determined by the architecture of production rather than by the policy rate. A deeply integrated economy faces a fatter right tail because its supply chains are long, and this imposes a structural constraint on what monetary policy can achieve.

4.1. **Environment**

Consider a set of countries $\{1, \ldots, M\}$. Each country m has the hub-chain production structure of Section 2, and all countries are exposed to the same global hub shock $\varepsilon > 0$. Goods markets are linked through standard open-economy demand channels: country m's IS schedule includes a real exchange rate term, denoted e_{m1}. Purchasing power parity holds in the long run. We assume that steady-state natural rates are identical across countries ($r_m^n = r^n$, $\forall m$), so any cross-country wedge in the natural rate reflects the hub shock alone.

Network exposure and aggregate supply. As in the closed-economy setting, the hub shock propagates through country m's production network so that the log-deviation of the price in sector j satisfies

$$\hat{p}_{mj} = \alpha_m \frac{1 - \delta_m^j}{1 - \delta_m} \varepsilon, \tag{45}$$

with the economy-wide aggregate exposure $\sum_j \omega_{mj} \hat{p_{mj}} = \Gamma_m \varepsilon$, where $\Gamma_m \equiv \Gamma(\alpha_m, \delta_m)$ is defined as in Equation (10). The country-m New Keynesian Phillips curve in period 1 is, therefore,

$$\pi_{m1} = \kappa_m y_{m1} + \beta \mathbb{E}_1[\pi_{m2}] + g_m \Gamma_m \varepsilon, \tag{46}$$

where $\kappa_m = g_m \omega_{mc,m}$ is the country-specific NKPC slope and the cost-push coefficient is $g_m \Gamma_m$, directly parallel to Equation (11). Natural output satisfies

$$y_{m1}^* = -\frac{\Gamma_m}{\omega_{mc,m}} \varepsilon. \tag{47}$$

The network multiplier Γ_m governs both the cost-push shift in the Phillips curve and the contraction in potential output simultaneously. The stagflationary dilemma of Section 2.3 is present in every country, but its magnitude is country-specific, determined by (α_m, δ_m) rather than by the shock itself.

World natural rate and country differences. From the closed-economy Equation (16), country m's hub-shocked natural real interest rate is

$$r^*_{m1} = r^n + \frac{\Gamma_m \varepsilon}{\zeta_m \omega_{mc,m}}. \tag{48}$$

In an open-economy equilibrium, saving and investment conditions are linked across borders. The world natural rate is, therefore, the expenditure-weighted average of country natural rates:

$$r_1^{n,W} = r^n + \varepsilon \sum_m \omega_m \frac{\Gamma_m}{\zeta_m \omega_{mc,m}}. \tag{49}$$

Country m's natural rate decomposes as world rate plus a country-specific wedge. Subtracting Equation (49) from Equation (48):

$$r^*_{m1} = r_1^{n,W} + \varepsilon \left[\frac{\Gamma_m}{\zeta_m \omega_{mc,m}} - \overline{\Gamma} \right], \tag{50}$$

where $\overline{\Gamma} = \sum_m \omega_m \frac{\Gamma_m}{\zeta_m \omega_{mc,m}}$ is the world-weighted average exposure.

The second term in brackets is country m's natural rate wedge: positive for economies whose network multipliers (adjusted for labour market and preference parameters) exceed the world average, negative for those below it.

Proposition 4. *A common global hub shock raises the world natural rate by $\sum_m \frac{\omega_m \Gamma_m \varepsilon}{\zeta_m \omega_{mc,m}}$. Countries with network multipliers above the expenditure-weighted world average experience positive natural rate wedges proportional to $\Gamma_m - \overline{\Gamma}$, where $\overline{\Gamma}$ is the appropriately weighted world average. These wedges scale linearly in Γ_m and are observationally distinct from the disinflation-cost premium: they enter the standard policy rule rather than augmenting it.*

Proof. Equation (50) follows directly from subtracting Equation (49) from Equation (48). Linearity in Γ_m is immediate from Equation (48). $\square$

Relative demand. Output gaps in country m satisfy the open-economy IS schedule

$$y_{m1} = -\zeta_m(i_{m1} - \mathbb{E}[\pi_{m2}] - r^*_{m1}) + \chi_m e_{m1}, \tag{51}$$

where e_{m1} is the real exchange rate and χ_m governs expenditure switching.[9] Since natural rates differ across countries from Equation (50), a common world real interest rate cannot simultaneously support all economies at potential. For any two countries H and F, the cross-country output gap differential is

$$y_{H1} - y_{F1} = -\zeta_H(i_{H1} - \mathbb{E}_1\pi_{H2}) + \zeta_F(i_{F1} - \mathbb{E}_1\pi_{F2}) + (\zeta_H r^*_{H1} - \zeta_F r^*_{F1}) + (\chi_H + \chi_F)e_{H1}. \tag{52}$$

The cross-country natural rate differential $\zeta_H r^*_{H1} - \zeta_F r^*_{F1}$ is proportional to the difference in network multipliers from Equation (50).

Inflation risk. Firms in country m combine global supply-chain signals with domestic private information. The precision of the common signal is governed by (α_m, δ_m) as in Equation (22), with $S_m = S(\alpha_m, \delta_m, \varepsilon)$ the domestic network stress measure. Period-1 inflation in country m is

$$\pi_{m1} = \mu_m \Gamma_m \varepsilon + u_m, \quad u_m \sim \mathcal{N}(0, \sigma^2_{u,m}), \tag{53}$$

[9] Note the exchange rate channel offers a partial offset. As H's currency appreciates against F's with diverging policy rates, the local currency burden of the hub shock falls in H and rises in F. This dampens cross-country cost-push dispersion without eliminating the asymmetry in $\mathcal{F}_m(S_m)$.

where μ_m is the average posterior belief about persistence in country m. The conditional variance of inflation is

$$\sum{}^2(S_m) = \Gamma_m^2(\alpha_m, \delta_m)\varepsilon^2\tau_0^2(S_m) + \sigma_{u,m}^2. \tag{54}$$

The common shock ε, therefore, generates cross-country heterogeneity in both the mean and the upper tails of the inflation distribution. Two countries sharing the same mean inflation outcome have conditional variances that differ by the factor $\left(\frac{\Gamma_H}{\Gamma_F}\right)^2$ whenever their network multipliers differ, so 90th-percentile inflation outcomes diverge more than proportionally relative to the median. This is the cross-country analogue of Proposition 2 and provides the structural basis for the quantile regression identification in Section 5.

Monetary policy divergence. Applying the analysis of Section 3.3 to country m, the optimal policy rule is

$$i_{m1}^* = r_1^{n,W} + \underbrace{\varepsilon\left[\frac{\Gamma_m}{\zeta_m\omega_{mc,m}} - \overline{\Gamma}\right]}_{\text{natural rate wedge}} + \underbrace{\overline{\pi}_m^N}_{\text{mean inflation}} + \underbrace{\frac{\kappa_m(1+\psi_{c,m})g_m\Gamma_m\varepsilon}{\zeta_m[\kappa_m^2(1+\psi_{c,m})+\lambda_m]}}_{\text{augmented cost-push response}}, \tag{55}$$

where

$$\overline{\pi}_m^N = \frac{\lambda_m G_m \Gamma_m \varepsilon}{\kappa_m^2(1+\psi_{c,m})+\lambda_m}.$$

The policy-rate differential between any two countries decomposes as

$$i_{H1}^* - i_{F1}^* = \underbrace{(r_{H1}^* - r_{F1}^*)}_{\propto\,\Gamma_H-\Gamma_F} + \underbrace{(\overline{\pi}_H^N - \overline{\pi}_F^N)}_{\propto\,\Gamma_H-\Gamma_F} + \underbrace{(\Delta_{i,H} - \Delta_{i,F})}_{\propto\,\Gamma_H-\Gamma_F} \tag{56}$$

where Δi_m is the policy gap from the country-m version of Equation (42). All three channels on the right-hand side are proportional to differences in network multipliers, and all scale linearly in

Γ_m. The policy differential is governed by supply-chain topology, but through a linear relationship across all three channels.

4.2. **Network Welfare Cost Divergence**

Divergence in network welfare costs. Each country faces the network welfare cost:

$$\mathcal{F}_m(S_m) = \frac{\psi_{c,m}}{2} \sum{}^2(S_m) = \frac{\psi_{c,m}}{2}[\Gamma_m^2 \varepsilon^2 \tau_0^2(S_m) + \sigma_{u,m}^2]. \tag{57}$$

The ratio of network welfare costs between any two countries is therefore

$$\frac{\mathcal{F}_{\mathcal{H}}(S_H)}{\mathcal{F}_{\mathcal{F}}(S_F)} \approx \frac{\psi_{c,H}}{\psi_{c,F}} \cdot \frac{\Gamma_H^2}{\Gamma_F^2} \cdot \frac{\tau_0^2(S_H)}{\tau_0^2(S_F)}, \tag{58}$$

when the demand-noise terms $\sigma_{u,m}^2$ are small. The cost divergence is the square of the multiplier ratio. Calibrating to national input-output tables with $\Gamma_{AUS} = 0.092$, $\Gamma_{NZL} = 0.098$, $\Gamma_{DEU} = 0.104$, $\Gamma_{KOR} = 0.105$, and $\Gamma_{SWE} = 0.117$, at peak network stress and holding $\psi_{c,m}$ constant across countries, $(\Gamma_{DEU}/\Gamma_{NZL})^2 = 1.13$, a 13% welfare premium that roughly doubles the gap separating their mean inflation outcomes. No combination of monetary policy responses can close this welfare gap: $\mathcal{F}_m(S_m)$ is invariant to $i_{1,m}$ by construction. The quadratic scaling of $\mathcal{F}_m(S_m)$ in Γ_m means that deeply integrated economies bear a structurally larger welfare burden from hub disruptions that lies permanently beyond the reach of monetary policy. Structural policies, such as supply-chain diversification or strategic inventory buffers, are necessary to complement monetary policy in deeply integrated economies.

Proposition 5. *Under a common global hub shock ε, optimal monetary policy differs across countries through three distinct channels, each scaling linearly in Γ_m, while welfare burdens diverge quadratically;*

(i). Natural-rate wedges. From Equation (50), country m's natural wedge is $\varepsilon\left[\frac{\Gamma_m}{\zeta_m\omega_{mc,m}}-\overline{\Gamma}\right]$, proportional to $\Gamma_m-\overline{\Gamma}$;

(ii). Mean inflation differentials. The optimal equilibrium inflation $\overline{\pi}_m^N=\frac{\lambda_m g_m\Gamma_m\varepsilon}{\kappa_m^2(1+\psi_{c,m})+\lambda_m}$ is proportional to Γ_m;

(iii). Disinflation-cost premia. The policy gap Δi_m from the country-m version of Equation (42) is proportional to $g_m\Gamma_m\varepsilon$, hence linear in Γ_m;

(iv). Network welfare cost divergence. The network welfare cost $\mathcal{F}_m(S_m)\propto\Gamma_m^2$ scales quadratically with network integration. Two countries with identical mean inflation outcomes but different network multipliers face welfare costs that differ by the square of their multiplier ratio.

Proof. (i). From Equation (50): the wedge contains Γ_m in the numerator and the wedge is proportional to $\Gamma_m-\overline{\Gamma}$; (ii). $\overline{\pi}_m^N$ contains Γ_m as its only country-specific factor in the numerator; linearity follows; (iii). Δi_m from the country-m version of Equation (42) contains $g_m\Gamma_m\varepsilon$ in the numerator; linearity follows (iv). $\mathcal{F}_m(S_m)=\frac{\psi_{c,m}}{2}\sum{}^2(S_m)$ contains Γ_m^2 from Equation (57). □

Figure 6a and 6b illustrate the policy gaps and network welfare costs for a range of countries. The policy gaps range from 20.5 basis points for the United States, which serves as the reference economy, to 26.3 basis points for Sweden, the economy with the highest hub intensity α from Table B.1. The pre-emptive policy gap is linearly related to the degree of network integration, Γ, implying that more integrated networks require stronger policy tightening to stabilise aggregate outcomes.

In contrast, the network welfare cost $\mathcal{F}\left(S\right)$ scales nonlinearly with network integration as Γ^2. As a result, welfare cost dispersion across economies is substantially larger than the dispersion in policy gaps. At the synchronisation levels, the model associates with a significant hub disruption, Sweden's network welfare cost is substantially more than Australia's, and is widening with S. The cross-economy dispersion in welfare costs is approximately double the dispersion in mean inflation outcomes.

Korea and New Zealand make for an instructive comparison precisely because their network architectures differ markedly. While Korea is an industrial dynamo, New Zealand has a strong agricultural focus. Yet their chain depths are similar in the calibration. What differs is the hub intensity, α. Korea's production network is characterised by fragmented manufacturing chains with significant energy and logistics embeddings, leading to larger welfare costs. In contrast, New Zealand's chain depth is driven primarily by agricultural processing, which lacks strong energy-sector exposure, resulting in a lower welfare exposure from the same shock.

Empirical predictions. The open-economy model generates four interconnected cross-country predictions that the empirical analysis in Section 5 exploits.

- **Quantile asymmetry in inflation outcomes.** Upper quantiles of the inflation distribution should respond more than proportionally to Γ_m relative to the median. Since $\sum^2(S_m) \propto \Gamma_m^2$ while $\overline{\pi}_m \propto \Gamma_m$, the $\Gamma_m^2{\cdot}\text{shock}_t$ interaction should dominate at quantiles $\tau \in (0.75,\, 0.90)$ while the linear $\Gamma_m{\cdot}\text{shock}_t$ interaction dominates at $\tau = 0.50$. This is a direct cross-country test of Proposition 2 and the quadratic scaling result.
- **Synchronisation amplification.** During synchronised hub shock episodes – when S_t is elevated – tail amplification should intensify in proportion to $\Gamma_m^2 \cdot S_t$, providing a separate identification of the belief-coordination channel distinct from the cost-push channel. Countries with deeper chains should experience disproportionate widening of their inflation fan charts relative to mean outcomes.

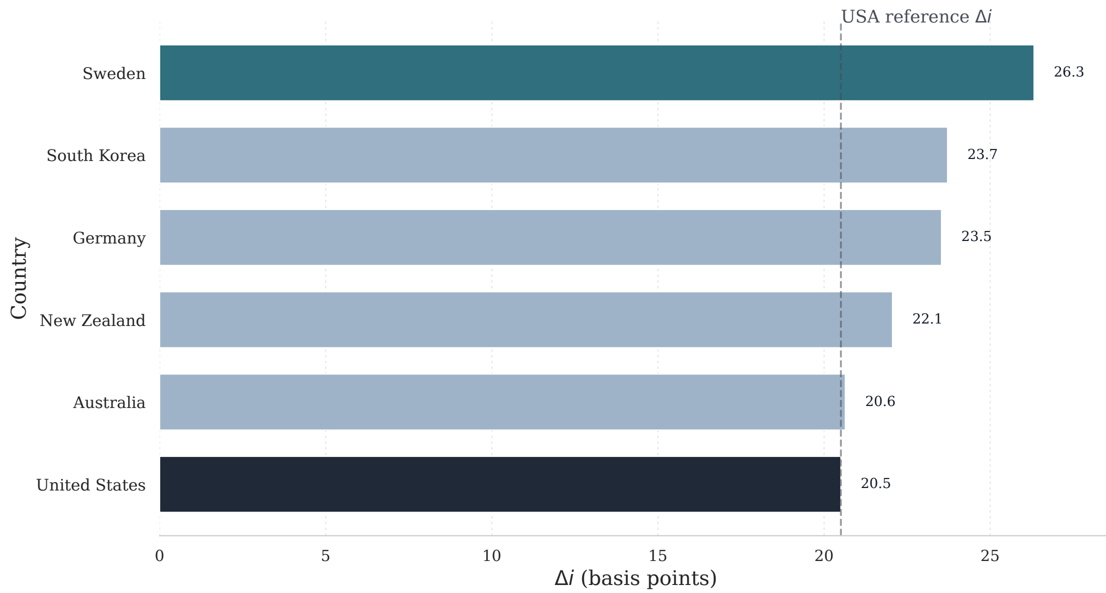

(a) Pre-emptive Policy Gap Δi (basis points above standard rule), calibrated with $\psi_c = 0.55$, $\kappa = 0.2$, $\zeta = 1$, $\lambda = 0.025$, $\epsilon = 0.5$, and $g = 0.116$. The United States sits at the dashed line (20.5 bps). Every other economy shown in the figure requires a larger pre-emptive tilt because its network multiplier, Γ, is higher, from Figure 3 and Table B.1.

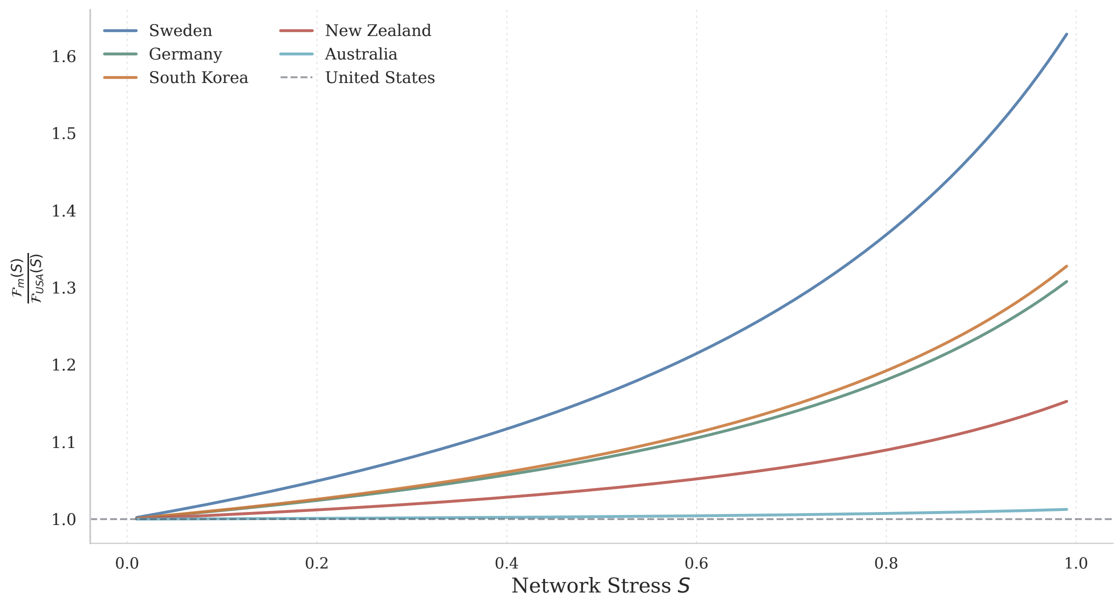

(b) Network Welfare Cost $\mathcal{F}(S)$ relative to the United States, $\frac{\mathcal{F}_m(S)}{\mathcal{F}_{USA}(S)}$, as a function of network stress S. The United States sits at the dash line ($\frac{\mathcal{F}_m(S)}{\mathcal{F}_{USA}(S)}$=1). Every other economy shows relatively higher network welfare cost than the United States, which grows in non-linearity with network stress S, from Equation (44).

FIGURE 6. Policy gaps and network welfare costs.

- **Linear interest rate divergence.** Countries with higher Γ_m should deviate more from global monetary conditions during supply disruptions, with the deviation linear in Γ_m. Panel regression of policy rate deviations from the world rate on $\Gamma_m \cdot \text{shock}_t$ should yield a positive coefficient. This tests channels (i)-(iii) of Proposition 5 jointly.
- **Quadratic network welfare cost divergence.** The network welfare cost $\mathcal{F}_m(S_m) \propto \Gamma_m^2$ implies that the unexplained component of macroeconomic welfare losses following hub shocks – the portion not accounted for by mean inflation or output outcomes – should be quadratic rather than linear in network exposure. In practice, this connects to the sacrifice ratio literature: at peak synchronisation, the output cost of restoring inflation to target should be disproportionately larger in deeply integrated economies, even after conditioning on their mean inflation overshoot. This provides a test of the network welfare cost that exploits its implication for the disinflation path without requiring direct measurement of $\mathcal{F}_m$.

Predictions 1–3 are directly testable using panel quantile regressions of inflation on supply-chain synchronisation measures interacted with country-level network exposure constructed from national input-output tables. Prediction 4 connects the framework to the sacrifice ratio evidence across disinflation episodes and provides an independent test of the quadratic mechanism.

5. Empirical Results

5.1. Data

We construct a monthly country panel to test if a common global hub disruption generates systematically different inflation-risk outcomes across economies with different production network architectures. The baseline sample covers 53 economies from January 2015 to December 2024, combining four sources of variation: realised inflation outcomes, predetermined country-level network exposure, global supply-chain stress, and exchange-rate movements. The merged panel contains 6,325 country-month observations.

Inflation data are drawn from the World Bank Global Database of Inflation. We use both headline and core CPI inflation at a monthly frequency as twelve-month rates. The monthly frequency preserves the timing of the February 2022 escalation of the Russia–Ukraine war, which we treat as the main hub-shock episode, and allows us to track the full conditional distribution of inflation during a period in which supply-chain stress, energy prices, and freight costs moved together. The dependent variable is realised inflation: we ask whether economies with deeper network exposure experienced larger shifts in the realised upper tail following the common shock.

Country-level network exposure is summarised by the network multiplier Γ_m, constructed from national input-output tables and predetermined relative to the shock episode. As derived in Section 2.2, Γ_m is the sufficient statistic for the mean cost-push effect of a hub disruption (equation 10), while Γ_m^2 governs the amplification of inflation risk through the belief-coordination channel (Proposition 2). Empirical specifications, therefore, include both Γ_m and Γ_m^2 interacted with common time-varying shock measures. Since Γ_m is time-invariant, country fixed effects absorb its level; identification comes from differential inflation responses to common shocks across economies with different ex ante exposure.

We proxy global supply-chain synchronisation using the Federal Reserve Bank of New York's Global Supply Chain Pressure Index (GSCPI). As established in Section 3.1, the relevant theoretical object is not the level of any single supply-chain indicator but the degree to which cost signals move together across sectors and countries, i.e. the network stress measure S, which equals the equilibrium belief-coordination weight ρ (Corollary 1). The GSCPI is an empirical proxy for S:

it captures the common component of supply-chain pressure across freight, manufacturing, and logistics indices in exactly the way that S captures cross-sector cost co-movement in the model. Since GSCPI is common across countries in any given month, its level is absorbed by time fixed effects. Empirical variation resides in the interactions $\Gamma_m \times GSCPI_t$ and $\Gamma_m^2 \times GSCPI_t$, which ask whether the same global synchronisation shock maps into different parts of the inflation distribution according to domestic network architecture.

Controls include nominal effective exchange rates from the BIS (absorbing conventional exchange-rate pass-through) and lagged inflation (accounting for persistence differences). Imposing data-availability requirements yields 6,034 country-month observations for headline inflation and 6,032 for core. A small number of extreme observations associated with hyperinflationary episodes were removed prior to estimation, retaining approximately 95 per cent of the merged panel.

5.2. **Empirical Methodology**

Baseline Differential-Exposure Design. The first empirical prediction, from Proposition 2, is that a common hub disruption generates disproportionate increases in the upper tail of inflation in economies with greater network exposure, with amplification scaling in Γ_m^2 rather than Γ_m. We test this using panel quantile regressions estimated at the median, the 75th percentile, and the 90th percentile of the conditional inflation distribution.

Let $\widetilde{\Gamma}_m = \Gamma_m - \overline{\Gamma}$, where $\overline{\Gamma}$ is the cross-country sample mean. Centring improves numerical stability and makes the linear interaction interpretable at the average-economy exposure level. Define the Russia–Ukraine shock indicator:

$$RU_t = 1, \quad \text{for } t \geq \text{February 2022}; \qquad RU_t = 0, \quad \text{otherwise.} \tag{59}$$

For inflation measure $k \in \{H, C\}$, corresponding to headline and core inflation, the baseline specification is:

$$Q_\tau\left(\pi^k_{m,t} \mid \mathcal{X}_{m,t}\right) = \alpha^k_{m,\tau} + \lambda^k_{t,\tau} + \beta^k_{1,\tau}\left(\widetilde{\Gamma}_m \times RU_t\right) + \beta^k_{2,\tau}\left(\widetilde{\Gamma}^2_m \times RU_t\right) + \rho^k_\tau \pi^k_{m,t-1} + \delta^k_\tau Dep_{m,t} + u^{k,\tau}_{m,t}, \tag{60}$$

where $\tau \in \{0.50, 0.75, 0.90\}$. Country fixed effects, $\alpha^k_{m,\tau}$, absorb time-invariant differences in inflation levels and network exposure, and month fixed effects, $\lambda^k_{t,\tau}$, absorb all shocks common to all economies. Identification comes entirely from differential inflation responses to the same global shock across economies with different predetermined network architectures. The coefficient $\beta^k_{1,\tau}$ measures the differential response attributable to the linear cost-push channel; $\beta^k_{2,\tau}$ measures the additional nonlinear amplification attributable to quadratic network exposure. The model yields two discriminating predictions: $\beta^k_{2,0.90} > \beta^k_{2,0.75} > \beta^k_{2,0.50}$, and the linear term $\beta^k_{1,\tau}$ should be near zero across quantiles once time fixed effects absorb the common mean cost-push. This second prediction distinguishes the belief-coordination mechanism from a purely mechanical cost-propagation story, which would generate a significant linear term and a flat quadratic.

Synchronisation Specification. The second prediction, from Corollary 1, is that upper-tail amplification should intensify when global supply-chain stress is synchronised. We augment the baseline with a triple interaction between quadratic network exposure, the shock episode, and standardised $GSCPI^z_t$:

$$\begin{aligned} Q_\tau\left(\pi^k_{m,t} \mid \mathcal{X}_{m,t}\right) =& \alpha^k_{m,\tau} + \lambda^k_{t,\tau} + \beta^k_{1,\tau}\left(\widetilde{\Gamma}_m \times RU_t\right) + \beta^k_{2,\tau}\left(\widetilde{\Gamma}^2_m \times RU_t\right) \\ &+ \beta^k_{3,\tau}\left(\widetilde{\Gamma}^2_m \times RU_t \times GSCPI^z_t\right) + \rho^k_\tau \pi^k_{m,t-1} + \delta^k_\tau Dep_{m,t} + u^{k,\tau}_{m,t}, \end{aligned} \tag{61}$$

where $GSCPI_t^z$ is the standardised index. Month fixed effects absorb both the average global inflationary effect and the common level of supply-chain pressure. The model predicts $\beta^k_{3,\tau} > 0$ with amplification strongest in the upper tail: $\beta^k_{3,0.90} > \beta^k_{2,0.75} > \beta^k_{3,0.50}$.

5.3. **Results**

Headline Inflation. Table 1 reports quantile estimates for headline inflation. Three features of results speak directly to the theoretical framework of Section 3.

Quadratic dominance in the upper tail. In columns (1)-(3) the coefficient on $\widetilde{\Gamma}^2_m \times RU_t$, is negative at the median (-25.80, significant at 1%), close to zero at the 75th percentile (-0.65, insignificant), and positive and highly significant at the 90th percentile (+46.67). The linear term $\widetilde{\Gamma}_m \times RU_t$ is near zero and imprecise at all quantiles. The presence of a zero linear term alongside a quadratic term significant only in the tail is the discriminating prediction of Proposition 2: pure cost pass-through would produce a significant linear term and a flat quadratic, whereas the belief-coordination mechanism generates the nonlinear tail amplification the data exhibit.

The negative median coefficient as corroboration of the welfare decomposition. The large negative coefficient on $\widetilde{\Gamma}^2_m \times RU_t$ at the median warrants careful interpretation. In the welfare decomposition of Section 3.3, expected losses split into a reducible component (the mean $\overline{\pi}^2$, amenable to interest rate policy) and an irreducible network welfare cost $\mathcal{F}(S)$ (governed by the variance $\sum{}^2(S)$, invariant to i_1). The augmented policy rule (Proposition 3) prescribes tighter policy (larger $\Delta i > 0$) in deeply integrated economies, and this tightening acts on the reducible mean component.

The negative median coefficient is consistent with this prediction: high-Γ_m economies that tightened more aggressively, in line with the augmented rule, show lower-than-average median inflation conditional on the common global shock and on month fixed effects. Meanwhile, the right tail remains elevated, reflecting $\mathcal{F}(S)$, i.e. the component that optimal monetary policy can internalise but cannot eliminate. The negative median coefficient alongside the positive 90th-percentile coefficient is the empirical counterpart of the welfare decomposition. The reducible component was partially addressed by tighter policy, but the structural network welfare cost remains in the tail.

Synchronisation as the activation channel. Columns (4)–(6) add the triple interaction $\widetilde{\Gamma}_m^2 \times RU_t \times GSCPI_t^z$. The coefficient $\beta_{3,\tau}^k$ is positive and precisely estimated across all three quantiles with magnitudes substantially larger at the upper percentiles – from 73.27 at the median, 109.94 at the 75th percentile, and 102.30 at the 90th percentile. The introduction of this term attenuates the negative median coefficient on $\widetilde{\Gamma}_m^2 \times RU_t$ (from -25.80 to -4.77), confirming that the median compression visible in columns (1)-(3) is partly a synchronisation interaction that operates differentially across the conditional distribution. The triple-interaction result is the empirical counterpart of Corollary 1. Synchronised supply-chain pressure increases the precision of the common signal v, tightening belief coordination and widening the upper tail in proportion to Γ_m^2.

Robustness: Core Inflation. Table 2 reports estimates for core inflation, which excludes food and energy components and provides a stricter test of propagation effects distinct from first-round commodity pass-through. The quadratic exposure coefficient rises monotonically across quantiles: 8.26 (significant at 10%) at the median, 23.21 (significant at 1%) at the 75th percentile, and 36.67 (significant at 1%) at the 90th percentile. Unlike the headline specification, the median coefficient is positive, consistent with the theory: core inflation is less directly exposed to the commodity-price pass-through that drives tighter monetary responses in deeply integrated economies, so the demand-suppression effect visible at the headline median is attenuated. The synchronisation interaction displays the same monotone ordering — 34.97, 49.33, and 68.37 at the three quantiles. The

TABLE 1. Headline Inflation: Quantile Asymmetry and Synchronisation

	Quantile asymmetry			Synchronisation amplification		
	(1) $\tau = 0.50$	(2) $\tau = 0.75$	(3) $\tau = 0.90$	(4) $\tau = 0.50$	(5) $\tau = 0.75$	(6) $\tau = 0.90$
$\widetilde{\Gamma}_m \times RU_t$	-0.071	0.045	1.810*	-0.259	0.129	0.036
	(0.415)	(0.587)	(0.988)	(0.415)	(0.585)	(0.992)
$\widetilde{\Gamma}^2_m \times RU_t$	-25.796***	-0.648	46.670***	-4.766	31.616***	66.341***
	(7.107)	(10.048)	(16.923)	(7.186)	(10.138)	(17.196)
$\widetilde{\Gamma}^2_m \times RU_t \times GSCPI^z_t$				73.269***	109.942***	102.299***
				(6.792)	(9.581)	(16.251)
Lagged headline inflation	0.988***	1.027***	1.060***	0.985***	1.023***	1.057***
	(0.001)	(0.002)	(0.003)	(0.001)	(0.002)	(0.004)
NEER depreciation	2.855***	2.890***	6.373***	2.872***	2.819***	6.126***
	(0.422)	(0.596)	(1.004)	(0.421)	(0.594)	(1.008)
Lagged inflation control	Yes	Yes	Yes	Yes	Yes	Yes
Exchange-rate control	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6034	6034	6034	6034	6034	6034
Pseudo ^{2}	0.836	0.867	0.889	0.837	0.869	0.890

concentration of both effects in core inflation confirms that the results are not driven by contemporaneous commodity pass-through but by deeper propagation through production networks and coordinated repricing beyond the most directly affected consumption components.

TABLE 2. Core Inflation: Network Exposure and Synchronisation

	(1) $\tau = 0.50$	(2) $\tau = 0.75$	(3) $\tau = 0.90$
$\widetilde{\Gamma}_m \times RU_t$	0.124	0.275	0.726
	(0.243)	(0.343)	(0.506)
$\widetilde{\Gamma}^2_m \times RU_t$	8.264*	23.212***	36.669***
	(4.223)	(5.948)	(8.779)
$\widetilde{\Gamma}^2_m \times RU_t \times GSCPI^z_t$	34.966***	49.333***	68.373***
	(3.959)	(5.577)	(8.231)
Lagged core inflation	0.990***	1.015***	1.032***
	(0.001)	(0.002)	(0.003)
NEER depreciation	0.324	0.843**	1.653***
	(0.256)	(0.361)	(0.533)
Lagged inflation control	Yes	Yes	Yes
Exchange-rate control	Yes	Yes	Yes
Observations	6032	6032	6032
Pseudo R^2	0.866	0.885	0.895

Robustness: Alternative Shock Window. Table 3 restricts the Russia–Ukraine shock period to February 2022–December 2023, focusing on the phase most directly associated with elevated energy, commodity, and supply-chain pressures. The main distributional pattern is preserved. The quadratic exposure coefficient at the 90th percentile rises sharply to 76.37, and the synchronisation interaction reaches 118.62 at the 75th percentile and 114.90 at the 90th percentile, both substantially above the median effect of 90.65. The ordering between the two upper quantiles is not strictly monotone, but both upper-tail coefficients far exceed the median, and the results are most pronounced precisely when the model predicts belief coordination to be most active, i.e. during the peak period of cross-sector cost synchronisation.

TABLE 3. Alternative Russia–Ukraine Shock Window

	Quantile asymmetry			Synchronisation amplification		
	(1) $\tau = 0.50$	(2) $\tau = 0.75$	(3) $\tau = 0.90$	(4) $\tau = 0.50$	(5) $\tau = 0.75$	(6) $\tau = 0.90$
$\widetilde{\Gamma}_m \times RU_t^{2022-23}$	-0.078	-0.811	0.858	0.227	-0.854	-1.866
	(0.510)	(0.693)	(1.196)	(0.502)	(0.719)	(1.169)
$\widetilde{\Gamma}_m^2 \times RU_t^{2022-23}$	-18.004**	20.172*	76.372***	-23.312***	28.384**	92.914***
	(8.626)	(11.731)	(20.237)	(8.513)	(12.189)	(19.816)
$\widetilde{\Gamma}_m^2 \times RU_t^{2022-23} \times GSCPI_t^z$				90.647***	118.618***	114.900***
				(7.181)	(10.282)	(16.715)
Lagged headline inflation	0.988***	1.026***	1.059***	0.986***	1.023***	1.056***
	(0.001)	(0.002)	(0.003)	(0.001)	(0.002)	(0.003)
NEER depreciation	2.914***	2.885***	6.197***	2.757***	2.826***	5.975***
	(0.423)	(0.575)	(0.991)	(0.416)	(0.596)	(0.969)
Lagged inflation control	Yes	Yes	Yes	Yes	Yes	Yes
Exchange-rate control	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6034	6034	6034	6034	6034	6034
Pseudo R^2	0.836	0.867	0.889	0.837	0.869	0.891

Four features of our results are worth drawing together.

Linear versus quadratic scaling. Proposition 2 predicts that the network-attributable component of inflation variance scales with Γ_m^2 while the mean cost-push scales with Γ_m. The consistent insignificance of the linear term $\beta_{1,\tau}^k$ across all specifications and quantiles, alongside the sharp significance of the quadratic term $\beta^k_{2,\tau}$ in the upper tail, directly supports this asymmetry. This is the key discriminating test between the belief-coordination mechanism and a simpler cost-propagation story: mechanical pass-through predicts a significant linear response with a zero quadratic, whereas coordination generates the nonlinear tail amplification the data exhibit.

Synchronisation as activation. The triple-interaction estimates confirm that synchronisation is the primary activation channel of the tail-widening mechanism, not merely a scaling robustness check. The $\beta^k_{3,\tau}$ coefficients dominate the $\beta^k_{2,\tau}$ main effect in the headline specification and rise sharply at the upper quantiles. This maps directly onto Corollary 1: it is the combination of deep network integration and synchronised cross-sector cost signals that triggers tight belief coordination and fattens the right tail. A central bank monitoring aggregate supply-chain indices without conditioning on their cross-sector synchronisation is observing a necessary but not sufficient condition for the mechanism to activate.

Welfare decomposition in the data. The joint pattern of a negative median and positive 90th-percentile $\beta^k_{2,\tau}$ in headline inflation can be read as a decomposition of the welfare split in Section 3.3. Deeply integrated economies that responded with tighter monetary policy compressed the median of their conditional inflation distribution. Yet their right tails remained elevated, reflecting the structural variance $\sum{}^2(S)$ invariant to i_1. This provides indirect evidence that the welfare decomposition is operative: optimal policy can internalise the reducible mean component but cannot eliminate $\mathcal{F}(S)$ in the tail.

The distributional nature of hub shocks. Across all specifications and both inflation measures, the quadratic and synchronisation effects are economically negligible at the median and large at the 90th percentile. A forecasting or policy framework calibrated exclusively to the conditional mean of inflation neglects dimension of risk the model identifies as structurally determined and largest in deeply integrated economies. The cross-economy dispersion in welfare costs from hub disruptions reflects this tail amplification. As such, it is a feature of the architecture of production and not of the shock itself. It also cannot be closed by any combination of monetary policy responses.

6. Conclusion

This paper identifies a structural asymmetry in how hub-chain production networks translate supply disruptions into inflation risk. The same network multiplier governs both the mean cost-push of a hub shock and the precision of the common signal that coordinates firm beliefs about its persistence. Since the variance of the inflation distribution scales with the square of the multiplier while the mean scales only linearly, two economies facing an identical shock with identical mean forecasts can face very different outcomes in the upper tail. Expected welfare costs decompose into two components: one amenable to monetary policy intervention and another, structural network welfare cost. Panel quantile evidence from the Russia-Ukraine conflict of 2021-22 episode is suggestive of a linear network-exposure term that is statistically indistinguishable from zero, alongside a quadratic term concentrated in the upper quintiles, which is the discriminating signature of belief coordination, not mechanical cost pass-through.

The optimal policy rule calls for greater monetary tightening than the standard prescription and provides a structural foundation for the risk-management tradition articulated by Greenspan (2004): when the costs of policy errors are asymmetric, optimal policy should tilt away from the baseline. In our model, the Greenspan principle is shaped by observable features of supply-chain architecture and does not require knowing how long the disruption will last.

Three practical implications follow. When a hub disruption occurs, the appropriate benchmark for rate-setting is the hub-shocked natural rate, not the pre-shock neutral – both the cost-push and the natural rate rise are governed by the same network multiplier, and both must be accommodated before anti-inflationary tightening can begin. The key monitoring object is the synchronisation of supply-chain stress indicators across sectors, not their level; an economy in which freight rates, energy costs, and delivery times are moving together faces a structurally different welfare problem from one in which only aggregate indices are elevated. This is the same synchronisation channel that dominates the empirical results in the upper quintiles, giving the monitoring prescription a direct empirical counterpart. And the time to respond is before the inflation distribution has fully widened, before belief coordination becomes self-reinforcing.

Beyond monetary policy, the returns to supply-chain diversification are convex – reducing hub intensity or chain depth generates welfare gains more than proportionally. It points to a role for structural measures like resilience investment and the creation of strategic inventories as a way of minimising the network welfare costs engendered by hub shocks.

References

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

Adrian, T., Grinberg, F., Liang, N., Malik, S., & Yu, J. (2022). The term structure of growth-at-risk. *American Economic Journal: Macroeconomics*, *14*(3), 283–323.

Angeletos, G.-M., & La'o, J. (2013). Sentiments. *Econometrica*, *81*(2), 739–779.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, *9*(4), 254–280.

Banerjee, R. N., Mehrotra, A., & Zampolli, F. (2020). *Inflation at risk from covid-19* (Technical Report No. 28). Bank for International Settlements.

Baqaee, D., & Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, *87*(4), 1155–1203.

Baqaee, D., & Farhi, E. (2022). Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. *American Economic Review*, *112*(5), 1397–1436.

Benigno, P., & Woodford, M. (2005). Inflation stabilization and welfare: The case of a distorted steady state. *Journal of the European Economic Association*, *3*(6), 1185–1236.

Boehm, C. E., Flaaen, A., & Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake. *Review of Economics and Statistics*, *101*(1), 60–75.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, *12*(3), 383–398.

Evans, C. (2019). Risk management and the credibility of monetary policy. *Speech at the New York Association of Business Economists Economic Luncheon, New York, New York, April*, *15*.

Greenspan, A. (2004). Risk and uncertainty in monetary policy. *American Economic Review*, *94*(2), 33–40.

Guerrieri, V., Lorenzoni, G., Straub, L., & Werning, I. (2022). Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages? *American Economic Review*, *112*(5), 1437–1474.

Lagarde, C. (2025). A robust strategy for a new era. *Speech at 25th "ECB and Its Watchers" Conference, Frankfurt, Germany*.

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

Ozhan, G. K., Sander, N., Wende, S., & Yang, S. (2025). *Monetary policy under network-level bottlenecks* (Tech. Rep.). Reserve Bank of Australia.

Rubbo, E. (2023). Networks, phillips curves, and monetary policy. *Econometrica*, *91*(4), 1417–1455.

Woodford, M. (2003). Optimal interest-rate smoothing. *The Review of Economic Studies*, *70*(4), 861–886.

Zhao, X., Duley, C., & Gai, P. (2026). *Local production architectures and global supply chain resilience: A network motif perspective.* University of Auckland, Mimeo.

Appendix A. Microfoundation for K

The common signal v is constructed from the CPI-weighted average of period 1 sectoral cost signals:

$$G \equiv \sum_j \omega_j s_j = \Gamma(\alpha, \delta)\varepsilon\theta + u, \quad u \sim \mathcal{N}(0, \sigma_s^2 \sum_j \omega_j^2). \tag{A.1}$$

Normalising $v \equiv G/(\Gamma\varepsilon) = \theta + u/(\Gamma\varepsilon)$. The precision of v as an estimator of θ is

$$\frac{1}{\sigma_v^2} = \frac{(\Gamma\varepsilon)^2}{\sigma_s^2 \sum_j \omega_j^2} = K\Gamma^2(\alpha, \delta)\varepsilon^2, \quad K \equiv \frac{N_{eff}}{\sigma_s^2}, \tag{A.2}$$

where $N_{eff} \equiv (\sum_j \omega_j^2)^{-1}$ is the effective number of sectors. The constant K depends only on the cross-sectional structure of the economy – CPI weights and idiosyncratic noise – not on the shock size. The shock ε enters through the loading $\Gamma\varepsilon$, separating structural features (captured by K) from cyclical ones (captured by ε). Coordination is strong when the economy is sectorally diverse (N_{eff} is large), supply-chain noise is low (σ_s^2 is small), and the shock is large (ε is large).

Appendix B. Supplementary Materials

B.1. Network Multipliers of Countries

Table B.1 shows the hub intensity, α, chain depth, δ, and network multiplier, Γ values across six economies.

Table B.1. Network Multipliers of Countries

	α	δ	$\Gamma(\alpha, \delta)$
Sweden	0.0788	0.3246	0.1166
Australia	0.0607	0.3366	0.0915
South Korea	0.0655	0.3766	0.1051
Germany	0.0708	0.3214	0.1043
New Zealand	0.0610	0.3756	0.0978
United States	0.0584	0.3570	0.0909