Shipping Lanes and Inflation-at-Risk:
Hub Shocks and Optimal Monetary Policy

Guest Lecture, NZ Treasury, 4 May 2026

Prasanna Gai

University of Auckland and Monetary Policy Committee, RBNZ

Preamble

Thank you, Struan, for that generous introduction and for the invitation to speak here today.

When I was starting my Honours thesis in Canberra, the United Kingdom had just joined the Exchange Rate Mechanism. The economics and politics of monetary union were live in a way that textbooks had not quite anticipated. I was provoked. I wrote an essay and gave it to my supervisor, the late Australian macroeconomist George Fane. He handed it back and said: go away and write a model.

Fast forward thirty-odd years. The Strait of Hormuz, a chokepoint which touches energy costs at almost every stage of global production, closes. I cannot help myself. I write another essay and show it to my colleague, the econometrician Peter Phillips. He hands it back and says: go away and write a model.

George was doing what great supervisors do. Peter had no such obligation. His message was simpler: the problem deserves serious analysis.

That is the approach this lecture takes. Models are not reality. But in a fast-moving episode, a simple model does something that careful watching and waiting does not: it tells you what you are actually looking at. And, in this case, what the standard framework has been missing.

I. The framework and its limits

Let me start with a thought experiment.

Take two economies. Hit them with the same global hub shock. Give them the same mean inflation forecast. The same central bank, the same policy rate response. Yet, one faces an output cost roughly double the other's. And no monetary policy response, however well-calibrated, can close that gap. The difference is not the shock. Not the policy. Not the forecast. It is the geometry of their supply chains.

So let me preview the main result upfront. The variance of the inflation distribution scales with the square of the network multiplier that governs cost pass-through, while the mean scales only linearly. Double the depth of your supply chains and you double the mean inflation problem. But you quadruple the tail risk and a disinflation cost that lies beyond the reach of monetary policy. The mean tells you where inflation lands. The network tells you how much it costs to have been wrong.

The rest of this lecture is about the logic behind these claims. But before the model, the policy context.

The monetary policy framework for handling supply shocks is well-founded. When an adverse cost disturbance hits the economy — a drought, for example — the first question is whether it is persistent. A temporary disruption raises the price level once; a permanent one shifts the inflation process. Tighten aggressively in response to a shock that would have resolved itself, and you inflict unnecessary output damage. Hold too low in the face of a permanent shock, and you allow pricing errors to embed in wage contracts and purchase a disinflation far more costly than the tightening you sought to avoid.

The ECB's 2025 strategy review states it clearly: assess persistence, look through what is temporary, respond only when second-round effects threaten to embed the disturbance. This is sensible and, in the context in which the textbook prescription was developed, close to the correct answer.

But Hormuz is not the context in which it was developed.

Two decades ago, Alan Greenspan articulated a guiding principle for risk-conscious central banking. When the costs of policy errors are asymmetric — when the damage from acting too little exceeds the damage from acting too much — optimal policy should tilt systematically toward the more consequential mistake. Not because the forecaster is wrong. But because a decision-maker who internalises the full distribution of outcomes will act differently from one who optimises only over the central case. The question the principle posed, but left open, is what determines whether the costs are asymmetric, and by how much. Applied today, the answer is written into the production architecture. And it is measurable.

What makes a hub shock different is not its size but its geometry. A peripheral shock originates at a point and travels, decaying as it goes, losing force with each step along the chain. A hub shock does not travel. It arrives. Everywhere, simultaneously, because every downstream sector draws from the same upstream source. That universality looks, at first, like what makes the problem complex. In fact, it makes the problem tractable.

When every firm receives the same signal at the same moment, something happens that the standard framework was not built to accommodate. Firms no longer form independent views about how long a shock lasts. They form views about each other's views. Persistence is no longer simply a property of the shock waiting to be revealed. It is partly constructed by the response to it — coordinated beliefs generate the repricing that validates them. A framework built for independent signals will systematically misread what is happening.

All the complexity, from the simultaneous arrival of correlated signals, the coordination of beliefs across thousands of firms, to the heterogeneity of exposure across economies whose supply chains differ in depth and intensity, can be captured by a single scalar, computable from standard input-output data, that I will call the *network multiplier*. It is, simultaneously, the cost-propagation mechanism, the belief-coordination device, and the determinant of which part of the welfare cost is within reach of the policy rate and which is not. Let me show why.

II. Production architecture

In recent work with Chanelle Duley and Cecilia Zhao, we show, using firm-level data across more than 100,000 firms and 190 countries, that global supply chains (Figure 1a) reduce overwhelmingly to two recurring structures: hub-spoke arrangements and outsourcing chains. The hub-chain is the empirically dominant form. What follows is its formalisation.

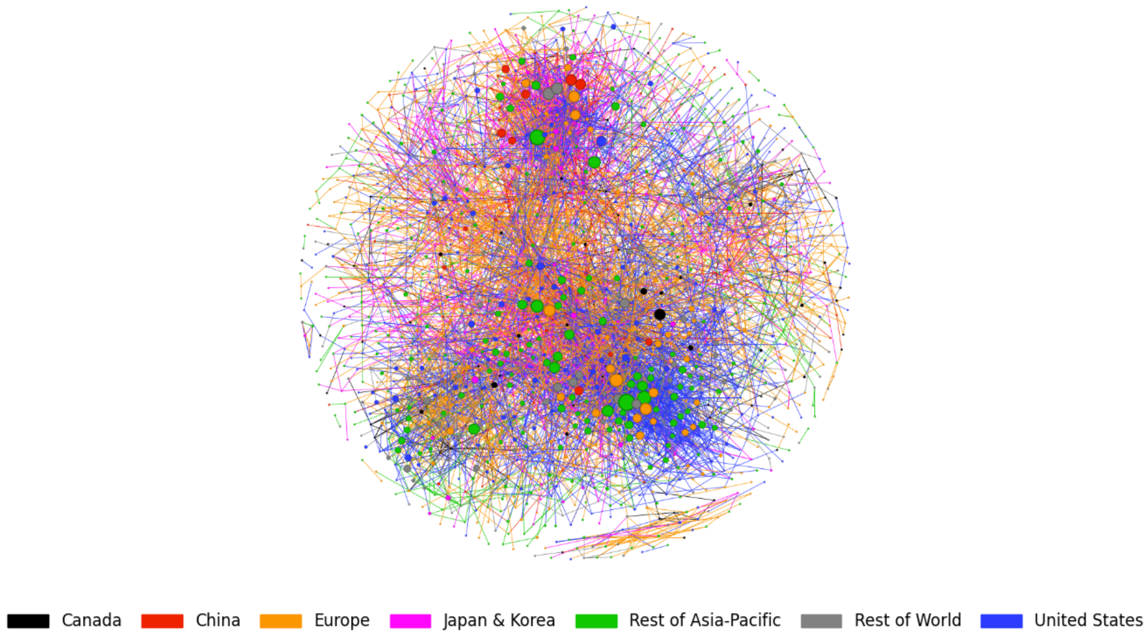

Figure 1a: The Global Supply Chain Network (July 2025)

At the top sits the hub — energy, logistics, the shipping lane — whose output is not one input among many but the indispensable ingredient in every sector's production, without exception.

Below it runs the chain: firms linked sequentially, each stage drawing on the one before it (Figure 1b). The question when the hub is struck is not which sectors are affected. It is whether any are not.

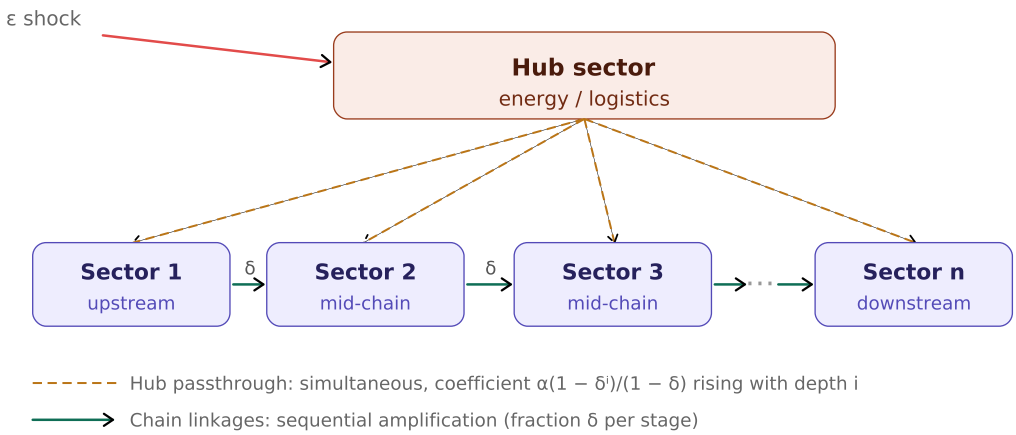

Figure 1b: Hub-Chain Production Network

When the hub is struck, the disturbance (ε) does not choose which sectors to visit. It arrives in all of them at once. The cost pass-through in sector i is:

$$\hat{p}_i \;=\; \frac{\alpha\left(1-\delta^{i}\right)}{1-\delta}\cdot\varepsilon \tag{1}$$

Two parameters carry everything. α (hub intensity) sets how deeply the hub is woven into each sector's costs. δ (chain depth) determines how much each sector inherits from the stage above. Together they tell you how far up the chain you feel the shock — and how little of it you escape.

III. The network multiplier

Weighting each sector's pass-through by its share of the consumer price index and summing across the economy gives the scalar that does the main work of the model:

$$\Gamma(\alpha,\delta) \;\equiv\; \frac{\alpha}{1-\delta}\sum_{j}\omega_j\,\left(1-\delta^{j}\right) \tag{2}$$

This is the *network multiplier* (Figure 2). α sets the floor; δ does the amplifying. Together they enter the Phillips curve:

$$\pi_1 \;=\; \kappa x_1 + \beta\,\mathbb{E}[\pi_2] + \kappa\,\Gamma(\alpha,\delta)\,\varepsilon \tag{3}$$

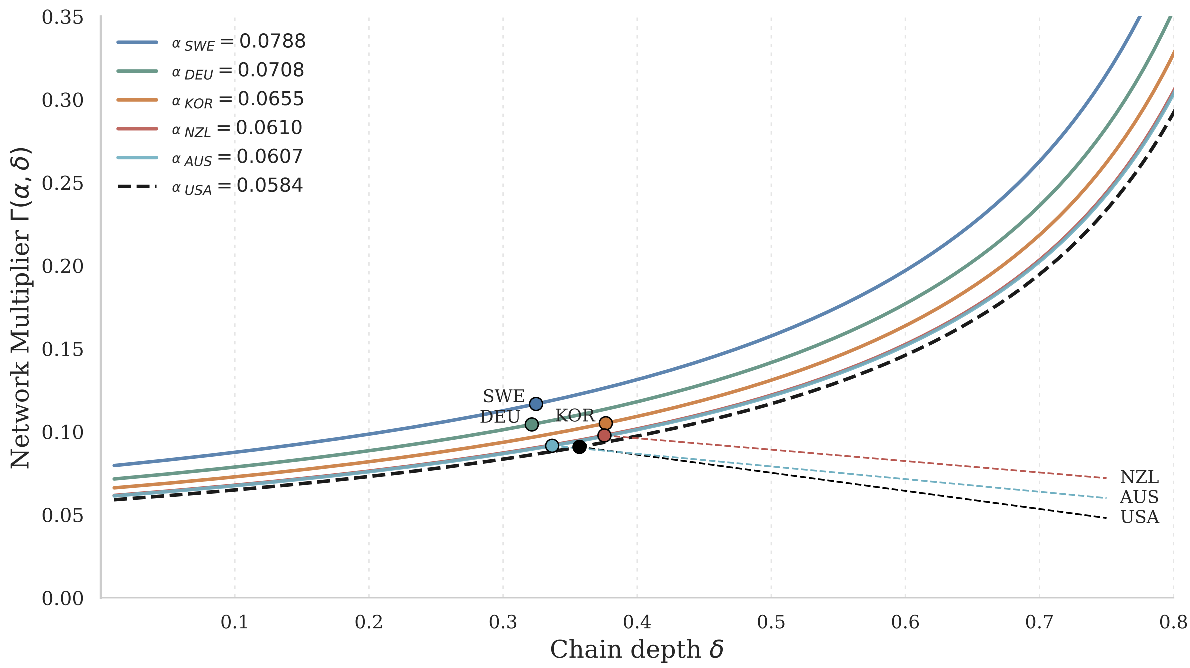

Figure 2: The Network Multiplier as a Function of Chain Depth

The term $\kappa\,\Gamma\varepsilon$ shifts the Phillips curve upward by an amount that has nothing to do with demand, nothing to do with the output gap, and nothing to do with firm behaviour. It is written into the production architecture before any pricing decision is made. Both α and δ come from standard input-output tables. All of the complexity of the supply chain compresses into two parameters and one number. And that same number governs the belief side of the model through exactly the same underlying geometry. The network multiplier is not just a cost-propagation device. It is simultaneously the source of the asymmetry stated in the introduction: variance scales with its square, and mean scales with the level.

IV. Horns of a dilemma

That number Γ is also, immediately, a problem for the central bank.

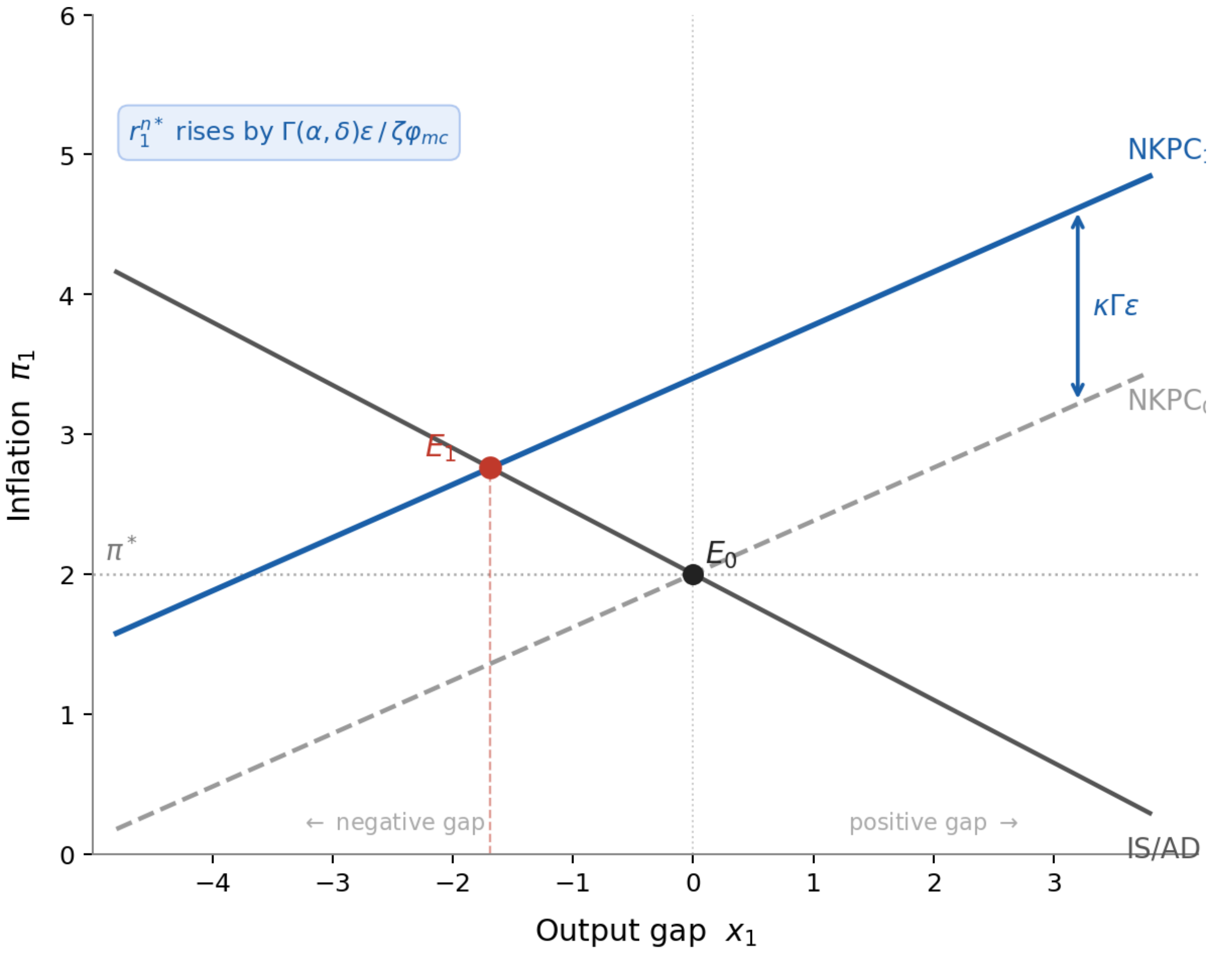

Figure 3: IS–NKPC Equilibrium

The hub shock does two things at once. It pushes the Phillips curve up by $\kappa\,\Gamma\varepsilon$ — inflation rises above target regardless of the output gap. And it raises the neutral rate of interest:

$$r_1^{n*} = r + \frac{\Gamma(\alpha,\delta)\,\varepsilon}{\zeta\cdot\varphi_{mc}} \tag{4}$$

Both move in proportion to Γ. ζ captures how readily households shift consumption across time; φ_{mc} captures how strongly real wages respond to the shock. Where labour supply is inelastic, costs are absorbed into prices rather than output and the inflationary impact is amplified. These parameters differ across economies, meaning that two countries with identical network multipliers can face different mixes of inflation and output loss from the same shock. Production architecture and labour market structure together determine the shape of the dilemma.

So the central bank faces a higher inflation rate and a higher neutral rate simultaneously. It must lift its policy rate just to stand still before it can begin to tighten. Raise rates beyond that, and you damage output that has already contracted. Hold rates low, and you validate the inflation. There is no free move. The stagflation is not incidental. It is the economy's own architecture, turned against it.

V. Correlated signals and the coordination of beliefs

The dilemma in Section IV is what the standard framework anticipates. What it does not factor in is the mechanism that determines how costly the inflation overshoot will prove to be.

When a hub is struck, every sector's cost signal contains the same common component – the hub shock amplified through that sector's position in the chain. Every sector draws from the same upstream source, so their cost experiences move together. This cross-sector co-movement is not noise. It is the identifying signature of a hub disruption and the structural precondition for what follows.

Each firm is trying to answer the same question: Is this permanent or temporary? It has its own private experience of rising costs, but it also observes the economy-wide co-movement of freight rates, energy prices, and delivery times. When those indicators move in lockstep, the common signal they carry swamps the firm's own private experience. Firms stop trusting their own judgement and start reading the room. The more synchronised the signals, the more their beliefs converge on a shared view. That degree of convergence is what ρ measures — the weight the average firm places on the common signal relative to its own private information:

$$\rho(\alpha, \delta, \varepsilon) = \frac{K\varepsilon^2\tau^2(\alpha, \delta)}{K\varepsilon^2\tau^2(\alpha, \delta) + \sigma_\eta^2} \tag{5}$$

where K measures the signal quality of the supply chain index, $\tau^2(\alpha, \delta)$ captures how strongly the network amplifies hub shocks into cross-sector cost covariance. And σ_η^2 is the variance of the idiosyncratic noise in each firm's private cost signal, the firm-specific component that is independent across sectors.

What determines ρ? Hub intensity α and chain depth δ — exactly the same geometry that determines Γ. The network does double duty: propagating costs on one side, coordinating beliefs on the other. A deeper, more hub-intensive economy is not just more exposed to the cost impact of a hub shock. It is simultaneously more susceptible to the belief entrenchment that makes that shock persist. And when beliefs coordinate on persistence, and the shock subsequently proves temporary, the pricing error must be unwound at an output cost whose expected magnitude grows non-linearly with network integration.

This is what economists have long called second-round effects. The conventional way to monitor them is to watch inflation expectations — surveys, market-implied measures and the like. These are valuable. But they are slow-moving by construction since they reflect beliefs that have already formed. By the time they signal that the inflation process is changing for the worse, the coordination has already occurred. The synchronisation of cost pressures across sectors is the leading indicator, not the lagging one.

It is natural to ask, if firms are coordinating more tightly and converging on the same view, might the inflation distribution get *wider*, not narrower?

The answer is that there are two distributions, and they move in opposite directions. When firms coordinate, they stop disagreeing with each other. That distribution narrows. But consider what coordination actually means here. Every firm is watching the same thing: freight rates up, port delays lengthening, energy costs rising. They each draw the same inference: this disruption looks permanent. They reprice accordingly. Nobody is holding out in the other direction to cancel the move.

If the Hormuz disruption does turn out to be permanent, no harm done and the pricing was right. But if it resolves and firms had all convinced themselves it would not, every one of them has made the same error simultaneously. There is no averaging out. The entire economy has mispriced together, and the central bank is left looking at an inflation overshoot it must now unwind across the board.

So the distribution of beliefs *across* firms gets tighter, reflecting agreement. But the distribution of aggregate inflation outcomes the central bank faces gets *wider*, because now everything hinges on whether the shared conviction was right. Coordination converts a thousand independent bets into a single one. And the central bank cannot see which way that bet has been placed.

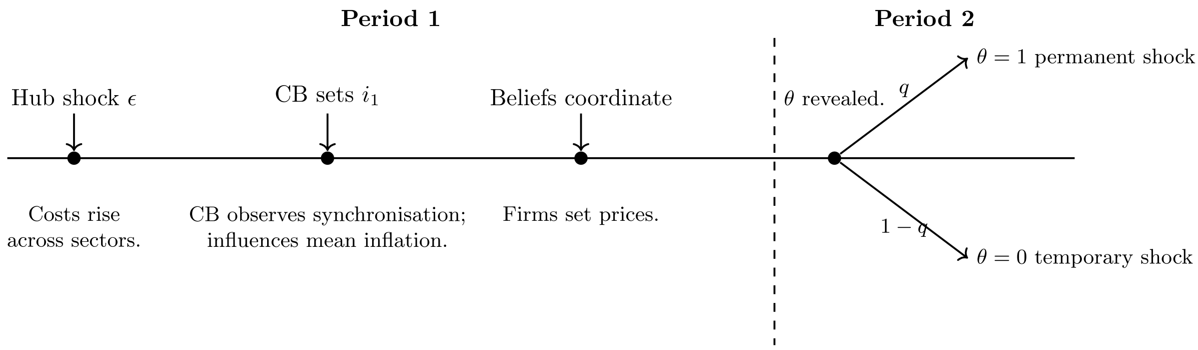

Figure 4: Model Timing — The Central Bank's Information Problem

The model timing makes the central bank's challenge precise. In period one, the hub shock arrives, costs move simultaneously, and the central bank must set the policy rate. The central bank observes S — the synchronisation of stress indicators — which tells it how tightly firm beliefs are likely to bunch. What it cannot observe is where those beliefs have settled. In period two, the shock's persistence is revealed. If permanent, period-one prices were right. If temporary, they were too high, and the error must be unwound. The central bank is not neglecting period two. It is correctly recognising that what happens there depends on something that lies beyond its reach at the time it must act.

VI. Network stress and the inflation distribution

There is one thing the central bank can observe that the model says matters most. Not where beliefs have settled — that is unobservable. Not how persistent the shock will prove — that is the central unknown. But whether the signals are moving together. That is S.

Define network stress as the share of cross-sector cost variation driven by the common hub source:

$$S \equiv \frac{K\varepsilon^2\tau^2(\alpha,\delta)}{K\varepsilon^2\tau^2(\alpha,\delta) + \sigma_\eta^2} \tag{6}$$

S is a signal-to-noise ratio in the most literal sense. When freight rates, energy prices, and delivery times all tell the same story simultaneously, S is high. When cost pressures are dispersed and sector-specific, S is low. A central bank monitoring S is, in effect, watching the coordination mechanism in real time, before it has done its damage.

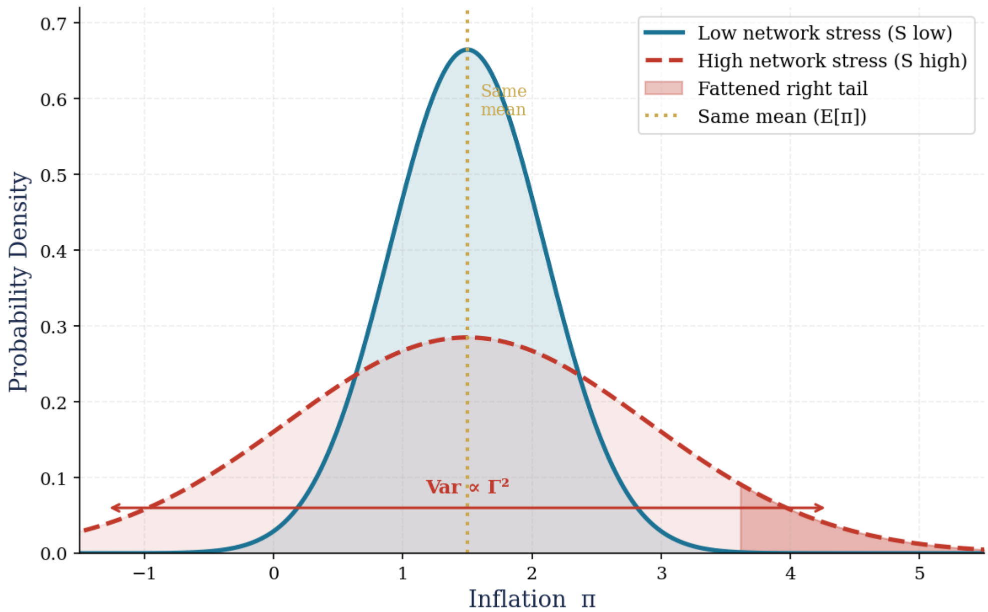

Figure 5: The Inflation Distribution at Low versus High Network Stress

Now here is where the claim I stated at the beginning of this lecture becomes an equation. The variance of inflation conditional on S is:

$$\Sigma^2(S) \;=\; \Gamma(\alpha,\delta)^2 \cdot \varepsilon^2 \cdot \tau_0^2 \cdot S \;+\; \sigma^2 \tag{7}$$

Variance scales with Γ^2. Mean scales with Γ. The same hub shock, the same mean forecast — but in a deeply integrated economy, the distribution is not just shifted, it is stretched. The right tail moves faster than the centre, and it moves with the square of the network multiplier, not the level. Two economies, one forecast, very different tails. The difference is determined entirely by where their Γ sits.

VII. Welfare decomposition and the optimal rule

Greenspan's principle has always been correct. It identifies the direction. But it does not specify the magnitude or the trigger. Asymmetric costs warrant pre-emptive action: yes. But toward what? By how much? Triggered by what? Let me give you a welfare decomposition that answers these questions and gives the asymmetry stated in Section I an analytical foundation.

The standard loss function is not wrong. It simply omits a structural term. Expected welfare loss in the model arises from two sources: the inflation and output-gap costs the central bank internalises in the standard way, and the additional output sacrifice required to unwind any period-one pricing error if the shock proves temporary. Putting both together, the total expected welfare loss decomposes as:

Total expected welfare loss = Monetary policy-responsive cost + Network welfare cost $\boldsymbol{\mathcal{F}(S)}$

The monetary policy-responsive cost is the familiar component — mean inflation and the output gap. Tighter policy reduces it. The tools work.

The network welfare cost $\mathcal{F}(S)$ is different. It is the expected disinflation burden attributable to the variance of the inflation distribution — the portion the central bank cannot reach:

$$\mathcal{F}(S) = \frac{\psi_c}{2}\Gamma^2(\alpha, \delta)\varepsilon^2\tau_0^2(S) \tag{8}$$

What determines it? The architecture of the supply chain, through Γ^2. The synchronisation of stress indicators, through S. Not the policy rate, which enters nowhere in equation (8).

So here is a precise statement of what the standard framework misses. The central bank observes S, which pins the width of the belief distribution. It cannot observe where beliefs have settled. It can compress mean inflation through the IS–NKPC channel. It cannot narrow the belief variance, because that is set by the network architecture and the shock. The timing of the model places $\mathcal{F}(S)$ beyond the reach of the central bank before the interest rate is even set. It is not a term the central bank forgets. But it is a term the model correctly locates outside its control. The friction in the model is thus not informational, but one of timing. By the time that the realisation of the common signal around which firms' pricing decisions have coordinated is revealed, the pricing error has already been made. In the language of economic theory, the central bank is constrained efficient. $\mathcal{F}(S)$ is what remains after optimal monetary policy has done everything it can.

What does reach $\mathcal{F}(S)$ is structural policy. Because the network welfare cost scales with Γ^2, anything that reduces hub intensity α or chain depth δ generates welfare gains more than proportionally. Strategic inventory buffers reduce the severity of the initial shock. Investment in alternative supply routes and supplier diversification lowers hub intensity directly. And supply chain resilience programmes that shorten or duplicate critical chains reduce δ. These are not precautionary luxuries. The returns to structural resilience are convex because the cost they reduce is quadratic.

A forward-looking central bank that internalises the period-two disinflation cost will nonetheless tighten more than the standard rule prescribes. Any period-one inflation above target must be unwound at an output cost, and a central bank that sees that future cost acts on it before it crystallises. The policy gap is:

$$\Delta i = i_N^* - i_Q^* \propto r\Gamma(\alpha, \delta)\varepsilon \tag{9}$$

The gap scales linearly with Γ and the shock ε. It does not require knowing how long the disruption will last, but only the geometry of the supply chain, available from standard input-output tables in real time.

So the welfare decomposition clarifies what Greenspan's principle left open. $\mathcal{F}(S)$ is the measure of the asymmetry. It is the structural welfare burden that grows quadratically with network integration that monetary policy cannot touch. Δi is the pre-emptive tilt: the reducible portion, correctly internalised by a forward-looking central bank. The asymmetry is structural. Pre-emption is not a posture. It follows from the network. And the trigger is observable before beliefs have coordinated and before the tail has fattened: the *synchronisation* of supply-chain stress across sectors, not its level.

VIII. Not all economies are alike

The opening claim of this lecture — two economies, one shock, one forecast, double the welfare cost — now has a precise source. *It is the ratio of their network multipliers, squared.*

The policy differential between any two countries decomposes cleanly:

$$\mathrm{i}_\mathrm{A}^* - \mathrm{i}_\mathrm{B}^* = \left(\mathrm{r}_\mathrm{A}^{\mathrm{n}^*} - \mathrm{r}_\mathrm{B}^{\mathrm{n}^*}\right) + (\pi_A - \pi_B) + \left(\Delta_{\mathrm{i,A}} - \Delta_{\mathrm{i,B}}\right) \tag{10}$$

The three channels: natural-rate wedge, mean inflation differential, disinflation-cost premium all scale linearly with Γ through the policy rate. What no rate decision closes is the divergence in network welfare costs, which scales quadratically. Different geometries, different welfare exposures, even after the optimal policy response.

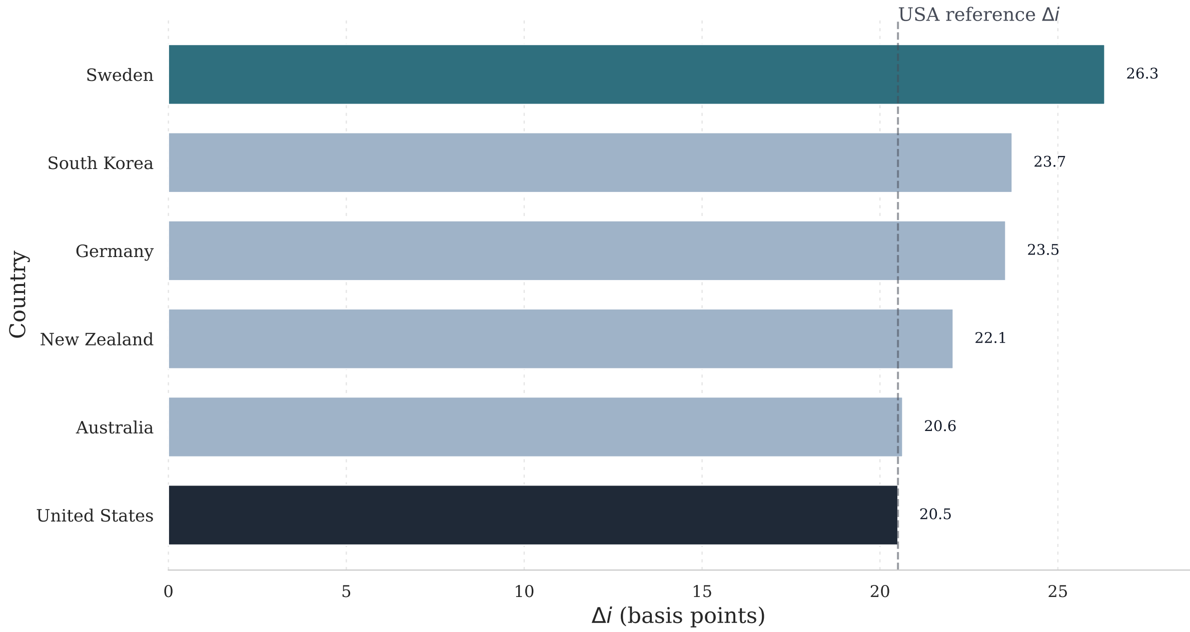

Figure 6a: Policy Gap — Cross-Country Comparison

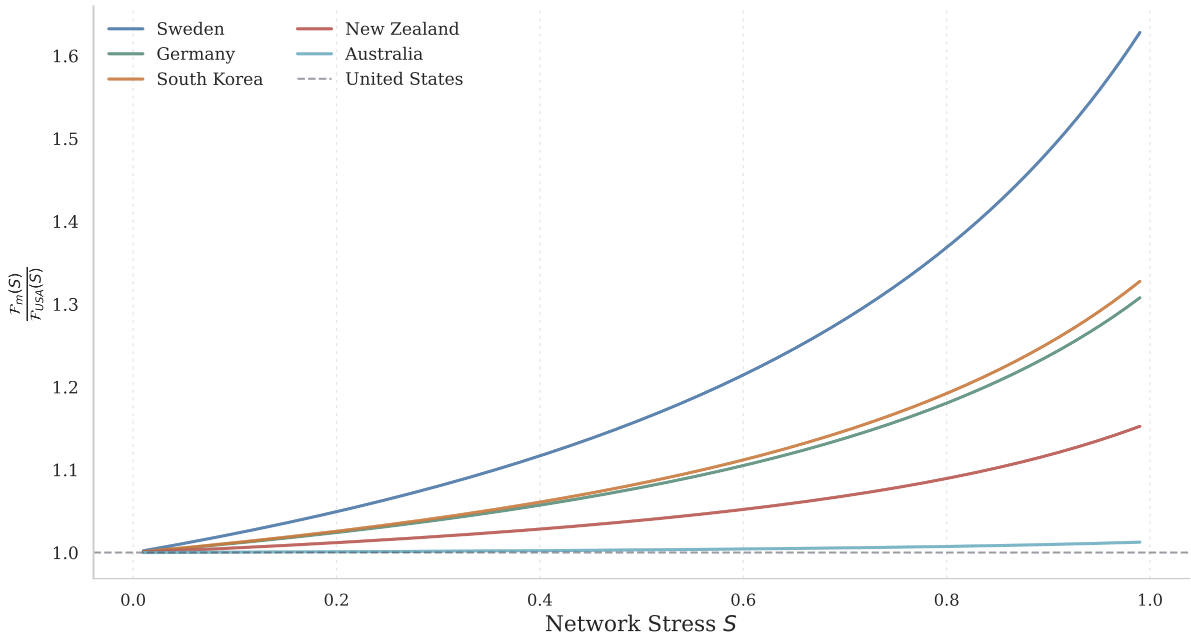

Figure 6b: Network Welfare Costs — Cross-Country Comparison

Some back-of-the envelope calculations help place the model in context. Calibrating to input-output data for five economies — Australia, New Zealand, Germany, Korea, and Sweden —

gives a sense of the magnitudes, though these should not be taken literally. Network multipliers range from around 0.09 for Australia to 0.12 for Sweden. Modest in levels. But because $\mathcal{F}(S)$ scales with Γ^2, the welfare dispersion is far larger. At the synchronisation levels the model associates with a significant hub disruption, Sweden's network welfare cost is substantially more than Australia's, and 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 they are so different. Korea is an industrial dynamo; New Zealand has a strong agricultural focus. The difference in welfare cost emerges from the fact that, in Korea, there is more "hub" woven into the base of each sector's costs than in New Zealand. This reflects its industrial structure – energy and logistics are more deeply embedded into heavy manufacturing and semiconductors.

IX. What the model implies

Three implications follow from the model, each one practically grounded.

First, when a hub disruption hits, the neutral rate has already risen before any decision is made as a structural consequence of the network topology. A central bank holding its rate constant may be easing, relative to a benchmark that has shifted.

Second, the synchronisation of stress indicators, not their level, is what warrants pre-emptive tightening. An elevated aggregate supply-chain index whose components are moving independently carries very different policy implications from one whose components are all pointing the same way simultaneously. The key diagnostic is not the height of the index. It is whether the components are telling the same story.

And third, waiting for inflation expectations surveys to confirm that second-round effects are underway is, in a hub-shock environment with high synchronisation, likely to mean waiting too long. By the time surveyed expectations have moved, coordination has already occurred and the disinflation cost is accumulating.

One misconception is worth addressing directly. Nothing here implies a reflexive tightening bias. As Willem Buiter reminded us when MPCs were first established, the monetary policy committee room is not an aviary. There are no hawks and no doves, only people trying to read what the evidence implies. Pre-emptive tightening is warranted only when synchronisation is high and the coordination mechanism is live. When cost pressures are dispersed and the inflation distribution is broadly symmetric, the model prescribes the standard response, nothing more. Peripheral disruptions, whose price effects decay rather than propagate universally, warrant the look-through the conventional framework has always recommended. Pre-emption is conditional on a structural signal. When that signal is absent, so is the case for it.

X. What Hormuz teaches us

The standard framework asked the right question: Is this shock persistent? What hub disruptions reveal is that it should be the second question, not the first. The first question is whether the economy's production architecture makes persistence the likely equilibrium, and whether the geometry of the supply chain is such that a common signal will coordinate beliefs before the shock has had time to be assessed independently. That is not a question about the disruption. It is a question about the network.

What this lecture has tried to show is that it can be answered, and from observable data. The network multiplier Γ compresses the production architecture into a single number. It governs the cost-push impact of the disruption. Through the same parameters, it governs the coordination of beliefs. And through the square of those same parameters, it governs the fattening of the inflation tail and the structural welfare cost that a mean-based framework cannot see. One number. Three roles. And each role follows from the same geometry.

Greenspan's risk management principle now has a structural answer. The asymmetry is Γ^2 versus Γ. The forecast is linear. The tail is quadratic. A central bank that reads only the mean is, in welfare terms, solving half the problem.

The model also clarifies who solves the other half. $\mathcal{F}(S)$ calls for different tools, perhaps strategic reserves, supply chain diversification, investment in alternative hub capacity. And a different institution. The returns, like the costs, are convex. That is a task for fiscal and structural policy.

I want to be honest about what the model does not do. It cannot dissolve the deeper uncertainty about whether the mechanism is operating in any particular episode. The framework clarifies the question. It does not eliminate the judgment.

A drought raises prices and passes. A hub shock raises prices, coordinates beliefs, and may embed itself in the inflation process through the very act of firms trying to assess whether it will. The difference between them is not one of degree. It is one of geometry. And geometry is something we can measure.

That is what Hormuz teaches us that a drought does not.

Thank you. If anyone would like a copy of the essay, it has now become a model.

** This lecture draws on joint work with Andy Haldane. The views expressed are my own and do not reflect those of the Reserve Bank of New Zealand or its Monetary Policy Committee. I thank Peter Phillips for challenging me to explore this problem, and I am grateful to Chris Bloor, Anna Breman, Caitlin Davies, and Hyun Shin for helpful comments and encouragement. Cecilia Zhao provided excellent research assistance. All errors are my own.*