

Firms in uncertain, dynamic networks: Unified strategy, risk and performance when perturbation theory breaks down

Authors: Dr. Tim De Jonckheere, Thibault Putseys

Executive summary. Firms operate inside networks of suppliers and customers, and in many economies a small "rich club" of large, densely interconnected firms forms the core of that network. Strategy, performance and risk management at firm level are nonetheless built on single-shock sensitivities set against a central scenario. This paper asks when that is enough. It is enough as long as three things hold: a shock is small relative to the network's distance to criticality, so that first-order perturbation theory is valid; the state of the network can be treated as known, although a firm in fact only holds a belief about it; and the network's couplings can be treated as settled, although in fact they are continuously reshaped by the actions of goal-directed firms. In that calm regime, sensitivities are sufficient statistics, and variance bridges, strategic sensitivity tables, value-at-risk and hedge ratios all work. In most circumstances, substitutes are scarce, shocks can push networks into a dynamic or even a critical regime, most of the time there is significant uncertainty about the state of the network and network couplings are dynamic. Then all three assumptions fail: the response becomes non-perturbative, uncertainty about the state dominates the tail, and the network's own dynamics take over. Superposition, attribution, sensitivity-based value-at-risk, risk contributions and plans built on fixed sensitivities then break down.

We show this on an illustrative network of 12 large and 150 small firms, and on a fully calculated five-firm network built from explicit income statements. In a calm state, a scenario of $+1$ percentage point inflation, $+1$ percentage point interest rates and a 10% fall in equity markets reduces the pre-tax profit of a small peripheral supplier by 25.6% and of a large hub manufacturer by 15.6%, within half a percentage point of the sum of their single-shock sensitivities. In a stressed state, the same scenario costs the supplier 118% of its pre-tax profit, breaching its liquidity buffer, and the hub 42%, against 41% and 20% implied by their sensitivities: the network's loop gain rises from 0.32 to 0.59 before the shock and reaches one during it. The 2022 European energy and inflation crisis shows the same signature in real data: German industrial production fell only 0.6% in 2022, but energy-intensive production was 19.6% lower by December and still 15.2% below its February 2022 level in March 2026, long after gas prices had normalised. A numerical simulation of that episode shows why a sensitivity-based method misses this persistence, while an active-inference-style simulation, in which the network's state and firms' own actions feed back into the outcome, captures it. We then draw out what this means for how scenarios are generated, introduce active inference as a framework that meets the requirements of the critical regime and reduces to sensitivity-based management in calm conditions, show how it translates into hedging risk, optimising performance and refining strategy, and introduce hedgeo, which applies these principles for European companies.

1. **Introduction:** Firms do not operate in a vacuum

Our earlier paper on European resilience [1] looked at the European economy as a single network: firms, sectors and countries connected by trading relationships, with a rich-club core of large, densely interconnected firms and institutions at its centre and the rest of the economy tied to that core, often indirectly. That paper asked what this structure means for the network as a whole – how shocks propagate across it, where concentration risk builds up, and how resilient the system is to a given disturbance.

Having established that the European economy itself has this shape, this paper turns the question around: given a network with a rich-club core, what does that structure mean for a single firm trading inside it? Every firm buys inputs from suppliers and sells outputs to customers, and each of those suppliers and customers has suppliers and customers of its own. Together they form a network of interlinked dependencies. The production-network literature has shown that this structure matters for the economy as a whole: shocks to individual firms or sectors travel along input–output links and can add up to aggregate fluctuations instead of averaging out [2,3,4]. A recent cross-country benchmarking of firm-to-firm datasets without reporting thresholds confirms these properties hold once incomplete data is controlled for, rather than being an artefact of which firms happen to be observed [5]. The evidence at firm level is direct, and recent. When a cyberattack shut down Jaguar Land Rover's UK production for over a month from the end of August 2025, the disruption did not stay inside JLR: many of its suppliers, dependent on JLR for the bulk of their revenue, began redundancy proceedings within weeks, an estimated 120,000 jobs across the supply chain were put at risk, and the UK government stepped in with a £1.5 billion loan guarantee [6]. Weeks later, a Dutch–Chinese dispute over the chipmaker Nexperia cut off roughly 80% of its chip output, and within days Honda had halved production at a North American plant and Ford and Volvo were warning of their own cuts, because the affected chips – simple components for locks, climate control and speedometers – had no readily available substitute [7]. Both episodes follow the pattern the production-network literature documents more systematically: after the 2011 Great East Japan Earthquake, disruptions spread to the customers and

suppliers of directly affected firms, and to their customers in turn [8]; natural disasters that hit a supplier reduce the sales of its customers, especially when the inputs are hard to replace [9].

The shape of the network also matters. Firm-level production networks have heavy-tailed degree distributions: a few firms have very many trading partners and most have few [10,11], as in scale-free networks [12]. Many complex networks also display *rich-club ordering*: the best-connected nodes are more densely connected to each other than chance would predict [13,14]. The resulting picture, a core of large interconnected firms and a large periphery of small firms tied to that core, is the set-up of this paper. Quantifying systemic risk directly on a nation-wide firm-to-firm supply network confirms this is not just a structural curiosity: a small set of firms carries a disproportionate share of the network's total systemic risk, concentrated in exactly the rich-club core [15].

Strategy, performance and risk management at firm level, by contrast, are all built for calm periods. Strategic plans are set against a central scenario; performance management tracks budgets and key performance indicators against that same baseline; and risk management is mostly built on sensitivities – interest-rate duration, equity beta, cost pass-through rates, profit-at-risk per risk factor – each measuring the response to one shock at a time, with the responses added up. This works when effects are small and additive, which is the domain of first-order perturbation theory, and when the environment can be summarised by a single baseline and a fixed vector of exposures around it. Six things break it in practice. The first is uncertainty about which regime the firm is actually in: a central scenario, or a sensitivity computed around it, says nothing about how likely a shock is, how costly it would be to have misjudged the regime, or when it is worth paying to find out. The second is model and parameter risk: a duration, a beta, or a profit-at-risk number is itself an estimate, calibrated on a historical window that may say little about the regime the firm is about to enter, so the sensitivity can be wrong before any shock ever arrives. The third is the network: Baqaee and Farhi show that, in a production network, the first-order approximation behind sensitivity-based management breaks down once shocks are large enough for second-order network effects to matter, and the result can then be badly wrong [16]; Elliott, Golub and Leduc show that supply networks can themselves be fragile, in the sense that small changes in the strength of relationships can cause discontinuous collapses of production [17]; modelling how a shock spreads through a firm's supply chain network, rather than assessing each counterparty's own financial condition in isolation, has been shown to raise standard risk measures by a factor of five or more, with the additional risk concentrated in a small set of firms [18]; and aggregating firm-level networks up to the sector level, the standard practice in macro models, throws away exactly the information needed to predict how a shock will propagate [19]. The fourth is observability: most firms have no visibility past their tier-one suppliers and customers, so the network that is missing from sensitivity-based management is often also missing from the firm's own data, regardless of how it is modelled. The fifth is organisational: budgets and key performance indicators built against a central scenario reward hitting that scenario, which discourages the people closest to the ground from flagging an early deviation from it. The sixth is latency, organisational and technological alike: even once a deviation is recognised, procurement contracts, approval cycles, and capital-allocation processes are built for calm-period cadence rather than for fast action, and the systems underneath them are usually no faster – data pipelines and reports that refresh overnight or monthly, and legacy platforms not built for a quick change of course, add their own delay on top of the organisation's. This last point is itself what active inference is built to confront, not a limitation of it: its case for acting early, while information is still cheap, only has force if the organisation can act on short notice, which is why treating the network as a whole has to be paired with more dynamic execution capabilities – policies that can be executed quickly once a belief about the regime has shifted, rather than waiting for the next planning cycle.

This paper takes the view of a single firm inside a rich-club network and asks when single-shock sensitivities are sufficient for strategy, performance and risk management, and what has to replace them when they are not. We answer the question for two firms that sit in different places in the same network: a small supplier in the periphery and a large manufacturer in the club. Section 2 sets up the network, the firm's linear response to macro-financial shocks and the condition under which perturbation theory is valid, and adds the two features that current practice leaves out: the state of the network is uncertain, so a firm holds a belief about it rather than knowing it, and firms act, so the network evolves over time as the outcome of their decisions. Section 3 shows that in a calm regime – shocks small relative to the network's distance to criticality, a concentrated belief and a settled action loop – sensitivities are sufficient for performance, strategy and risk alike, using a scenario in which inflation and interest rates rise while equity markets fall; it then shows how all three break down in the critical regime, where the response becomes non-perturbative, uncertainty about the state dominates and the network's own dynamics take over, including what this does to value-at-risk and risk contributions, and closes with the four requirements this sets for any framework that is to work in that regime. Section 4 draws out what these requirements mean for how scenarios are generated, Section 5 introduces active inference as a framework for measuring impact under network effects that meets them, and Section 6 shows how to act on it, to hedge risk, optimise performance and refine strategy. Section 7 introduces hedgeo and Section 8 concludes. Appendix A collects the notation, Appendix B derives the expansion parameter, and Appendix C specifies the two models used. Appendix D works the argument of Section 3 numerically, on an illustrative network with 12 large and 150 small firms, and Appendix E repeats it on a fully calculated example with explicit income statements, for a small supplier in the periphery and a large manufacturer in the club, and applies the same yardstick to the 2022 European energy crisis, including a numerical simulation comparing perturbation theory with active inference on that episode. All numbers in Appendix E are calculated from its own data (Tables 2–4); the code that produces them and every figure of the paper is published with it, and a separate calculation memo documents each step.

2. **Set-up:** A rich-club network of firms

This section and the next develop the general argument on an illustrative network with many small firms. Appendix E then repeats the analysis on a fully calculated example with explicit income statements.

2.1 The network

Consider n firms, split into a rich club $\mathcal{H}$ of m large firms and a periphery $\mathcal{P}$ of $n-m$ small firms, with $m \ll n$. Dependencies are collected in a matrix $A = (a_{ij})$, where $a_{ij} \geq 0$ is the share of firm j's activity that relies on firm i, either because i supplies j (an input share, as in input–output models [20]) or because i buys from j (a sales share). The column sum $\sum_i a_{ij} < 1$ is the total external dependence of firm j. Ordering the firms with the club first gives the block form

$$A = \begin{pmatrix} A_{\mathcal{H}\mathcal{H}} & A_{\mathcal{H}\mathcal{P}} \\ A_{\mathcal{P}\mathcal{H}} & A_{\mathcal{P}\mathcal{P}} \end{pmatrix}, \qquad \phi(k) = \frac{2E_{>k}}{N_{>k}(N_{>k}-1)}. \tag{1}$$

$A_{\mathcal{H}\mathcal{H}}$ is dense: large firms trade heavily with each other. $A_{\mathcal{P}\mathcal{P}}$ is sparse: small firms have few links among themselves. $A_{\mathcal{H}\mathcal{P}}$ and $A_{\mathcal{P}\mathcal{H}}$ hold the links between the periphery and the club: most small firms supply or buy from one to three large firms. The rich-club coefficient $\phi(k)$ is the density of links $E_{>k}$ among the $N_{>k}$ firms with more than k links; rich-club ordering means that $\phi(k)$ rises with k faster than in a randomised network with the same degrees [14]. Figure 1 shows such a network.

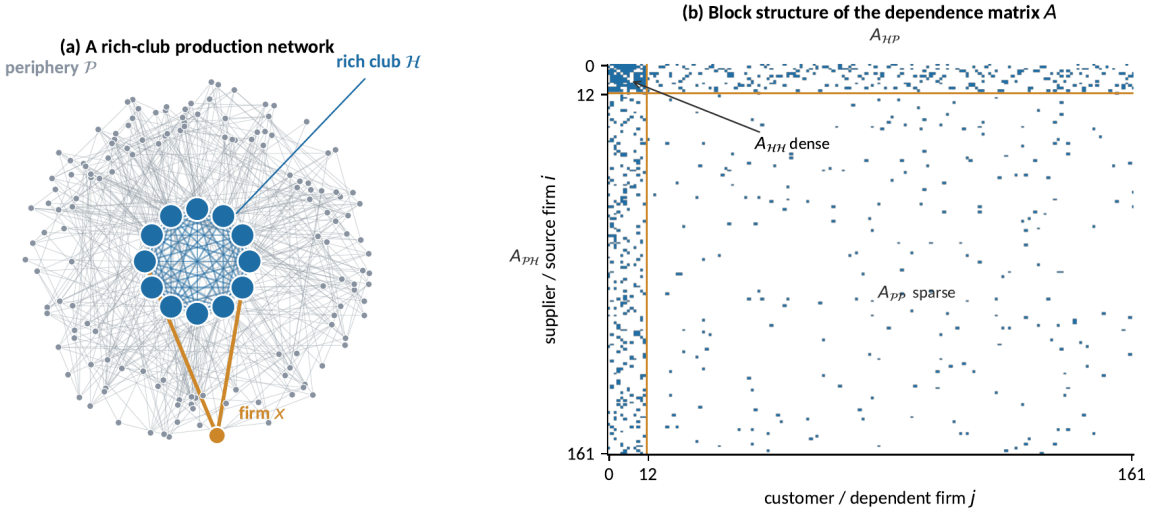


Figure 1. (a) Illustrative rich-club production network with 12 large firms (club $\mathcal{H}$) and 150 small firms (periphery $\mathcal{P}$). Firm x is a small firm linked to two large firms. (b) Non-zero entries of the dependence matrix A: a dense club block, sparse periphery block, and periphery–club links. Illustrative network of Appendix C.1, generated in code.

Not every dependence is equally binding. Following [1], write

$$a_{ij}^{\text{eff}}(\theta) = a_{ij}\,[1 - s_{ij}(\theta)], \qquad s_{ij} \in [0,1], \tag{2}$$

where a_{ij} is structural dependence, set by technology, contracts and infrastructure, and $s_{ij}(\theta)$ is the share of that dependence that can be substituted in state θ: alternative suppliers, other customers, spare capacity, inventory. The *loop gain* of the network is the spectral radius $\rho_{\text{eff}} = \rho(A^{\text{eff}})$: the factor by which a disturbance is multiplied each time it goes around the network.

2.2 Firm x, shocks and performance

Let $\ell \in \mathbb{R}^n$ be the vector of activity losses of all firms relative to a baseline (ℓ_i is firm i's shortfall in sales, output or cash flow, as a share of its baseline level), and $z \in \mathbb{R}^K$ a vector of macro-financial variables. In the example below $z = (\pi, r, e)$: inflation, interest rates and the fall in equity markets. Firm i has a direct exposure $\xi_i(z)$ to the macro shock, a loss that would occur even if i had no suppliers or customers (its own funding cost, its own final demand). On top of that, firm j absorbs a share a_{ij}^{eff} of every one of its counterparties' losses, so, firm by firm,

$$\ell_j = \xi_j(z) + \sum_i a_{ij}^{\text{eff}}\,\ell_i,$$

or, stacking all n firms into one vector equation,

$$\ell = F(\ell, z) := \xi(z) + A^{\text{eff}\top}\ell. \tag{3}$$

The transpose is there because a_{ij}^{eff} is indexed on j's dependence on i (it sits in column j of A^{eff}), while Eq. (3) must sum, for each firm j, the losses of every i that j depends on – and that sum is row j of $A^{\text{eff}\top}$, as the component form above makes explicit. Eq. (3) defines ℓ implicitly, as a *fixed point* of F: it is solved by finding the loss vector that reproduces itself once it has gone once around the network, not by evaluating a formula. We write $A^{\text{eff}}(\ell, z)$, rather than treating it as fixed, because substitutability depends not only on the state θ that Eq. (2) left open but also on ℓ and z themselves: substitutes become scarce when many firms are already stressed, and financing a switch of supplier becomes harder when rates rise. Where the state the network starts from needs to be explicit, we write $F(\ell, z; \theta)$. This is what allows the coupling to change as a shock unfolds, which Section 3 shows is the source of everything that goes wrong with sensitivity-based risk management. Firms also have buffers (liquidity, inventory, covenant headroom); when a loss exceeds a firm's buffer, the firm stops performing and its loss jumps. Firm x denotes the firm whose risk we track through the paper; Appendix E repeats the whole argument for a small firm in the periphery and, side by side, for a large hub. Its performance indicator $y_x = g_x(\ell_x, z)$ can be EBIT, free cash flow or equity value.

2.3 Linear response and the role of the network

Around a baseline (ℓ^*, z_0), write $J = \partial F/\partial \ell$ and $B = \partial F/\partial z$. A small shock δz gives

$$\delta \ell = (I - J)^{-1} B\, \delta z = \sum_{k=0}^{\infty} J^k B\, \delta z, \qquad \rho(J) < 1. \tag{4}$$

The k-th term is the effect arriving after k steps through the network: directly ($k = 0$), through suppliers and customers ($k = 1$), through their suppliers and customers ($k = 2$), and so on. For firm x,

$$\delta y_x = \Delta_x^\top \delta z + \tfrac{1}{2}\, \delta z^\top \Gamma_x\, \delta z + O(\|\delta z\|^3), \qquad \Delta_x = \partial_z g_x + \partial_\ell g_x\, e_x^\top (I - J)^{-1} B. \tag{5}$$

The sensitivity vector Δ_x is a network object: it contains firm x's direct exposure and everything that reaches x through the chain. In the rich-club set-up, most of that indirect exposure runs through the club, so it is natural to ask what the club looks like once the periphery's reaction to a shock has been folded in.

To see this, partition $\delta \ell = (\delta \ell_{\mathcal{H}}, \delta \ell_{\mathcal{P}})$ and write out the linear system $(I - J)\delta \ell = B \delta z$ of Eq. (4) block by block, club rows and periphery rows separately. The periphery row reads $\delta \ell_{\mathcal{P}} = (I - J_{\mathcal{P}\mathcal{P}})^{-1}\big(B_{\mathcal{P}} \delta z + J_{\mathcal{P}\mathcal{H}} \delta \ell_{\mathcal{H}}\big)$: it says how the many small firms respond once the hubs have moved, including the loops the periphery makes among itself, which is why $(I - J_{\mathcal{P}\mathcal{P}})^{-1}$ appears. Substituting this back into the club row eliminates $\delta \ell_{\mathcal{P}}$ altogether and leaves an equation in $\delta \ell_{\mathcal{H}}$ alone,

$$\big(I - S_{\mathcal{H}}\big)\, \delta \ell_{\mathcal{H}} = B_{\mathcal{H}} \delta z + J_{\mathcal{H}\mathcal{P}} (I - J_{\mathcal{P}\mathcal{P}})^{-1} B_{\mathcal{P}} \delta z, \qquad S_{\mathcal{H}} = J_{\mathcal{H}\mathcal{H}} + J_{\mathcal{H}\mathcal{P}}\, (I - J_{\mathcal{P}\mathcal{P}})^{-1} J_{\mathcal{P}\mathcal{H}}. \tag{6}$$

This elimination step – solving a block of a linear system for one set of variables and substituting it back – always produces a *Schur complement*, here $S_{\mathcal{H}}$, and it is standard even outside networks, for example in eliminating nuisance parameters from a linear regression. $S_{\mathcal{H}}$ is the effective coupling among the m hub firms once every possible route through the periphery has been accounted for: the first term, $J_{\mathcal{H}\mathcal{H}}$, is the direct hub-to-hub coupling, and the second term adds every loop that leaves a hub, is passed through one or more small firms – possibly circulating among them many times, which is exactly what the Neumann series $(I - J_{\mathcal{P}\mathcal{P}})^{-1} = \sum_k J_{\mathcal{P}\mathcal{P}}^k$ sums up – and eventually returns to a hub. In words: even though no single small firm matters much on its own, the periphery matters collectively, because $S_{\mathcal{H}}$ adds up its contribution over all $n - m$ small firms and every path length. Because the club block is dense and the spectral radius of a non-negative matrix is at least that of any of its diagonal blocks [21], the network's largest eigenvalue, and therefore its closeness to criticality, is set mainly by $S_{\mathcal{H}}$, that is, mainly in the club. A small firm feels this through its links to large firms, although it cannot observe it directly.

2.4 The network is uncertain

So far the state θ that sets effective dependence through $a_{ij}^{\text{eff}}(\theta) = a_{ij}[1 - s_{ij}(\theta)]$ has been treated as if it were known. In practice it is not: no firm, however well-informed, observes directly how much slack its suppliers' suppliers have, how tight credit is three tiers upstream, or how close the network as a whole sits to its critical point. What a firm has instead is a belief – a probability distribution $q(\theta)$ over the space Θ of possible states, which need not coincide with the true, unobserved law $\mathbb{P}(\theta)$ the world is actually drawn from. The network a firm should reason about is therefore not a single matrix computed at an assumed θ, but a random object,

$$A^{\text{eff}} = A^{\text{eff}}(\theta), \qquad \theta \sim q(\theta),$$

and every quantity built on it inherits that randomness: the loop gain $\rho_{\text{eff}}(\theta) = \rho\big(A^{\text{eff}}(\theta)\big)$, the loss vector $\ell(\theta, z)$ solving Eq. (3), and the distance to criticality $1 - \rho_0(\theta)$ of Eq. (7) are all random variables carrying a distribution induced by $q(\theta)$, not

single numbers a firm can simply look up. A firm that treats θ as known – plugging in a single, assumed state and computing a single ρ_{eff} – is implicitly working with a degenerate belief, $q(\theta) = \delta_{\theta_0}$, a point mass on one value. Section 3 shows when that simplification is harmless and when it is not.

2.5 Firms act: network dynamics

The network is not a fixed backdrop either: it is continuously reproduced, period by period, by the actions firms take. Let $d_i \in \mathcal{D}_i$ denote the decision firm i takes in a given period – how much to order, what price to charge, how much inventory to hold, which supplier to favour, how large a hedge to put on – and write the profile of every firm's decision as $d = (d_1, \ldots, d_n)$. Two things then happen, and they happen together. First, the state θ that the next period starts from is set not only by this period's losses and the next shock but by what every firm has just chosen:

$$\theta_{t+1} = \theta(\ell_t, z_{t+1}, d_t), \qquad a^{\text{eff}}_{ij,t+1} = a_{ij}\big[1 - s_{ij}(\theta_{t+1})\big],$$

so a firm that hoards inventory, locks in a supplier or tightens its payment terms changes $s_{ij}(\theta_{t+1})$, and hence A^{eff}_{t+1}, for every other firm that shares that supplier or that counterparty. Second, each firm chooses its action to serve its own objective Π_i, given what it expects the state to be and what it expects every other firm to do:

$$d_{i,t} = \arg\max_{d_i \in \mathcal{D}_i} \mathbb{E}_{\theta_t \sim q_{i,t}(\theta)}\Big[\Pi_i\big(d_i, d_{-i,t}, \theta_t, z_t\big)\Big],$$

where $q_{i,t}(\theta)$ is firm i's own belief about the state at time t (Section 5 describes how this belief is updated) and $d_{-i,t}$ is what it expects every other firm to do. These two equations close a loop: beliefs and expectations about others shape what each firm does; the state θ_t sets A^{eff}_t and, through Eq. (3), ℓ_t, for the period; and the resulting ℓ_t, what every firm has just done and the next shock z_{t+1} together set θ_{t+1} and the beliefs that open the following period. The network is therefore better described as the state of a dynamical system, $\big(A^{\text{eff}}_t, \ell_t, d_t, q_t\big)_{t \geq 0}$, than as a fixed graph that a shock is passed through once. Section 3 shows when this loop settles into a fixed point quickly enough to be ignored, and when it does not.

2.6 When is perturbation theory valid?

The expansion (5) is useful when the higher-order terms are small. The subtle point is that J is not fixed: it depends on the state and on the shock. To first order, $J(z_0 + \delta z) \approx J_0 + J'\delta z$, and a short calculation (Appendix B) shows that the series converges when

$$\kappa := \frac{c\,\|\delta z\|}{1 - \rho_0} < 1, \qquad \rho_0 = \rho(J_0), \tag{7}$$

where c measures how strongly the shock erodes substitutability. The ratio of second-order to first-order terms is of order κ. A shock is therefore not small or large in absolute terms; it is small or large relative to the network's *distance to criticality* $1 - \rho_0$. In a calm state, $1 - \rho_0$ is large and ordinary shocks give $\kappa \ll 1$. In a stressed state, $1 - \rho_0$ is small and the same shock can give κ close to one. This analysis holds a single, known state θ_0 fixed, and a network whose couplings have already settled – the degenerate case of the belief $q(\theta)$ and the action loop of Section 2.5 just described. What follows makes precise when that double simplification is harmless, and what fails when it is not.

3. **Regime change:** Where perturbation theory breaks down, uncertainty bites and dynamics take over

3.1 The calm regime: static, certain and settled

Current practice in risk, performance and strategy management rests on exactly the two simplifications just described: it treats the network's state as known rather than as a belief $q(\theta)$ held over a range of possibilities, and it treats the network's couplings A^{eff} as settled rather than as the live outcome of every firm's ongoing decisions d_t. Both are harmless in a calm regime, and this section makes precise why. Suppose firm x faces K shocks at once, collected in $\delta z = \sum_k \varepsilon_k e_k$ where e_k is the k-th unit vector (a pure inflation shock, a pure rate shock, and so on) and ε_k its size, and suppose $\kappa \ll 1$, so that Eq. (5) is dominated by its first (linear) term. Everything in this subsection is then a direct consequence of one fact: to leading order, the map from shocks to performance, $\delta z \mapsto \delta y_x \approx \Delta_x^{\top} \delta z$, is *linear*, with a single fixed sensitivity vector Δ_x that does not depend on which shock is being asked about or how large it is. Two further simplifications make this tractable: the belief $q(\theta)$ is concentrated enough – a point mass, or close to one, on a single state θ_0 – that $\mathbb{E}_{\theta \sim q}[\,\cdot\,] \approx (\cdot)\big|_{\theta_0}$, so a single Δ_x computed at θ_0 stands in for the whole distribution; and the action loop of Section 2.5 has already settled into a fixed point d^* that moves

slowly enough, relative to the shocks being analysed, that $A_t^{\text{eff}} \approx A^{\text{eff}}(\theta_0)$ can be treated as constant over the period in question – a *quasi-equilibrium*. What follows from linearity, a concentrated belief and a settled action loop together splits naturally into three domains – performance, strategy and risk – each correct up to terms of relative size $O(\kappa)$, and each is exactly the property a manager relies on when reporting performance, setting a plan, or running a hedge off single-factor sensitivities.

Performance

Superposition. The response to a combined scenario is the sum of the responses to each shock taken on its own:

$$\delta y_x\Big(\sum_k \varepsilon_k e_k\Big) = \sum_k \Delta_{x,k}\,\varepsilon_k\,\big[1 + O(\kappa)\big], \tag{8}$$

where $\Delta_{x,k} = e_k^\top \Delta_x$ is the k-th entry of Δ_x: firm x's sensitivity to shock k alone. This is just linearity of Eq. (5) in δz; it holds because a linear map applied to a sum of vectors equals the sum of the map applied to each vector, and the curvature term Γ_x that would spoil this is $O(\kappa)$ smaller than the linear term (Appendix B).

Attribution. Superposition can be read backwards: given a realised, already-happened combination of shocks, the resulting change in performance can be split back into the part caused by each driver, $\delta y_x = \sum_k \Delta_{x,k}\varepsilon_k + \mathrm{Res}_x$ with $\mathrm{Res}_x = O(\kappa)$ a small residual. A risk report that says "60% of this quarter's EBIT miss came from inflation, 30% from rates and 10% from equities" is making exactly this claim, and it is a reliable claim only while $\kappa \ll 1$. In an FP&A team, this is literally the logic behind a variance bridge: last quarter's EBIT miss is split into "X from FX, Y from volume, Z from cost inflation" using sensitivities computed once – a cost pass-through rate, a volume elasticity – and carried forward unchanged from one reporting cycle to the next. The bridge is trustworthy not because of accounting convention, but because the firm's κ is small enough that Superposition and Attribution actually hold. Both properties also lean on the two simplifications above: Δ_x is evaluated at a single assumed state θ_0 rather than averaged over a belief $q(\theta)$, and it is evaluated on a network whose couplings have already settled at the action loop's fixed point d^*, so that y_x can be treated as a function of z alone rather than of z and a still-moving d_t.

Strategy

A strategic choice – how much capacity to hold, which supplier mix to commit to, how large a hedge to put on – can be written as a decision d_x that firm x sets to maximise its objective of Section 2.5, written $\Pi_x(d_x, z)$ here because the state is held at θ_0 and every other firm's decision at its settled value d^*_{-x}, with y_x of Eq. (5) simply Π_x evaluated at whichever d_x the firm has chosen. The optimal choice $d_x^*(z) = \arg\max_{d_x} \Pi_x(d_x, z)$ satisfies the first-order condition $\partial_{d_x}\Pi_x(d_x^*, z) = 0$, and differentiating this condition in z – the implicit function theorem – gives

$$\Delta_x^d := \frac{\mathrm{d}d_x^*}{\mathrm{d}z} = -\big(\partial^2_{d_x d_x}\Pi_x\big)^{-1}\partial^2_{d_x z}\Pi_x,$$

the *strategic sensitivity*, a strategic analogue of the sensitivity vector Δ_x. Δ_x^d is a fixed number exactly when the curvature terms $\partial^2_{d_x d_x}\Pi_x$ and $\partial^2_{d_x z}\Pi_x$ do not themselves move with the shock – the same $\kappa \ll 1$ condition as before, since the only route by which they could move is the network resolvent $(I - J)^{-1}$ inside Π_x. To leading order, then, $d_x^*(z_0 + \delta z) \approx d_x^*(z_0) + \Delta_x^d\,\delta z$: the optimal strategic choice moves linearly and predictably with the environment. This is precisely why a strategic plan can be set once against a central scenario and revisited only occasionally – its own sensitivity Δ_x^d, like Δ_x, is a fixed number computed once rather than a problem to be resolved from scratch every time the environment shifts. In an annual planning cycle, Δ_x^d is what a strategic sensitivity table is standing in for: "if oil moves \$5, add one more rig", "if churn moves 2pp, revise the hiring plan by N heads" are statements about a fixed Δ_x^d, computed once during planning and used to nudge the plan as forecasts move, rather than triggering a full re-plan every time. Δ_x^d inherits the same two simplifications: it is computed by differentiating Π_x at a single assumed θ_0 rather than under a belief $q(\theta)$, and it presumes every other firm's decision d_{-x} has already settled at d^*_{-x}, so that x's first-order condition $\partial_{d_x}\Pi_x = 0$ can be solved holding everyone else's action fixed, rather than as one equation in the simultaneous system of Section 2.5.

Risk

Aggregation. If the shocks have covariance Σ_z (how much they move together, not just how large each one is on its own), then, because a linear combination of the components of a random vector has variance $\Delta_x^\top \Sigma_z \Delta_x$,

$$\sigma^2(\delta y_x) = \Delta_x^\top \Sigma_z \Delta_x, \qquad \mathrm{VaR}_\alpha \approx z_\alpha\,\sigma(\delta y_x), \qquad \mathrm{RC}_k = \frac{\Delta_{x,k}(\Sigma_z\Delta_x)_k}{\sigma(\delta y_x)}. \tag{9}$$

Here VaR_α, the value-at-risk at confidence level α, is the loss that is exceeded with probability $1 - \alpha$ (for example $\alpha = 99\%$), and z_α is the corresponding quantile of the standard normal distribution; this approximation treats δy_x as normally

distributed, which is reasonable when it is a fixed linear combination of shocks that are themselves close to normal. RC_k, the risk contribution of shock k, is the share of the total variance attributable to that shock once its correlation with the others is taken into account; because $\Delta_x^\top \Sigma_z \Delta_x = \sum_k \Delta_{x,k}(\Sigma_z \Delta_x)_k$ by construction, the RC_k add up exactly to $\sigma(\delta y_x)$, a property risk teams call *full allocation*. Every dependence between shocks enters only through Σ_z: the sensitivities themselves stay fixed.

Hedging. Because Δ_x does not move with the scenario, a hedge instrument with its own sensitivity vector h can be sized once, by minimising the variance of the hedged position, $w^* = \arg\min_w (\Delta_x - wh)^\top \Sigma_z (\Delta_x - wh)$, and the resulting hedge ratio w^* stays valid as the scenario unfolds – it does not need to be recomputed shock by shock or revisited as the shock grows. This is the mathematics behind a treasury or risk function's standard toolkit: a limit expressed in duration or beta, a VaR number produced from a covariance matrix re-estimated once a quarter, and a hedge ratio set at inception and left alone until the next rebalancing date. None of that machinery needs revisiting shock by shock, which is exactly why it can run on a quarterly cycle rather than being recomputed for every piece of news. The covariance Σ_z and the hedge ratio w^* rest on the same pair of simplifications: a single assumed state θ_0, so Σ_z need not itself be averaged over $q(\theta)$, and a settled action loop, so Σ_z need not account for how other firms' own hedging and ordering decisions might shift in response to the same shock.

In this regime, single-shock sensitivities are sufficient statistics for performance, strategy and risk management alike: everything the firm needs to budget, attribute, plan and hedge is contained in the sensitivity vector Δ_x and its strategic counterpart Δ_x^d, together with the shock covariance Σ_z, which can be estimated once and reused across every scenario the firm cares about, rather than resolving the whole network model (3) for every combination of shocks it wants to examine. This does not mean that the network is irrelevant: the network is inside Δ_x through $(I - J)^{-1}$, and by Eq. (6) that resolvent is set mainly by the club. It means that once Δ_x and Δ_x^d are known, the network need not be revisited to answer performance, strategy or risk questions – until, as Section 3.2 shows, the state of the network changes enough that neither is a fixed number to rely on any longer.

3.2 The critical regime: non-perturbative, uncertain and dynamic

Keep the same firm, the same network and the same family of shocks, and change only the state θ that Eq. (2) left open. Buffers have been drawn down by earlier shocks, alternative suppliers are scarcer and get taken up faster as system-wide stress rises, and credit is tighter, so that a rise in rates makes switching supplier harder exactly when switching is most needed. None of this changes the firm, the network or the shock; it changes $\rho_0 = \rho(J_0)$, and with it the distance to criticality $1 - \rho_0$ in Eq. (7). A scenario that gave $\kappa \ll 1$ in the calm state can give κ close to one in the stressed state, even though it is the same scenario applied to the same firm. Figure 2 shows the typical shape of what follows on the illustrative network (Appendix D works out the numbers): for small scenario intensities λ the full model and the linear expansion agree, then the full model bends away, jumps to a high-stress branch once substitutability collapses – the hysteresis described in [1] – and drops further once peripheral firms breach their buffers. Sensitivities measured in the stressed state are already larger than in the calm state, which is a first warning, but a larger sensitivity is still a single fixed number, and the point of this section is that no fixed number describes what follows. The same two simplifications that made Section 3.1 tractable now work against the firm. First, when $1 - \rho_0(\theta)$ is itself close to zero for a non-trivial part of the belief $q(\theta)$'s support, the point-mass approximation $\mathbb{E}_{\theta\sim q}[\,\cdot\,] \approx (\cdot)\big|_{\theta_0}$ breaks down exactly where it matters most: a small chance of a much worse state dominates the expectation rather than being averaged away. Second, the action loop of Section 2.5 need not settle: when hoarding and pre-emptive ordering feed on themselves, the fixed point d^* either does not exist, is not unique, or is not reached before the next shock arrives, so A_t^{eff} cannot be treated as the constant backdrop Section 3.1 assumed. Four mechanisms are at work: (a) amplification near criticality, (b) state-dependent coupling, (c) thresholds and (d) dependence between the shocks themselves. Each is taken up below where it bites, under performance, strategy or risk.

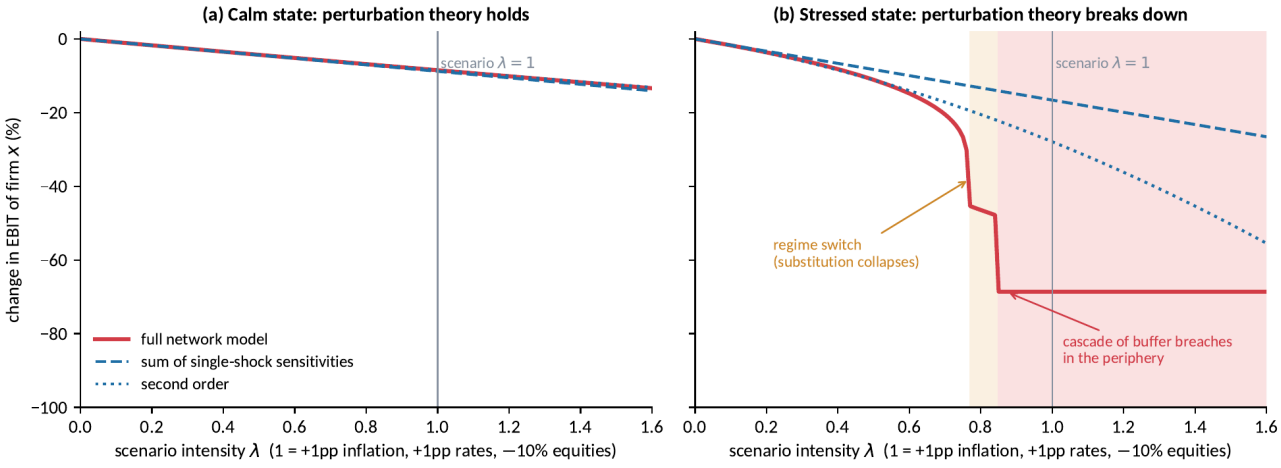


Figure 2. Change in EBIT of firm x as a function of scenario intensity λ. (a) Calm state: full model, sum of single-shock sensitivities and second-order expansion nearly coincide. (b) Stressed state: the full model departs from the expansion, jumps to a high-stress branch when substitutability collapses (regime switch) and then drops again when periphery firms breach their buffers. Illustrative network of Appendix C.1.

Performance

(b) State-dependent coupling. Because J depends on the shock, the expansion parameter κ of Eq. (7) approaches one and the series in Eq. (4) stops converging. In the scalar version of Appendix B, the cross-derivative between two shocks is

$$\frac{\partial^2 y}{\partial \varepsilon_1 \partial \varepsilon_2} = \frac{\bar{\xi}_1 c_2 + \bar{\xi}_2 c_1}{(1-\rho_0)^2}, \tag{10}$$

so an interaction term that is negligible next to the first-order terms in a calm state (small $\bar{\xi}_1 c_2$, large $1-\rho_0$) grows faster than those first-order terms as $\rho_0 \to 1$: the two shocks stop being separable. Appendix D quantifies this on the illustrative network. This is exactly where the performance toolkit of Section 3.1 breaks: superposition fails because the interaction term above is no longer negligible next to the first-order terms, and attribution fails because the residual Res_x in $\delta y_x = \sum_k \Delta_{x,k}\varepsilon_k + \mathrm{Res}_x$ is no longer small. It is the same size as, and can dominate the contributions a performance report attributes to each driver. In practice this is the moment where the quarter's EBIT no longer decomposes cleanly into the usual drivers, the "unexplained" or "other" line in the variance bridge balloons, and restating last quarter's attribution with this quarter's sensitivities gives a different answer from the one reported at the time. **Uncertainty.** Once κ is no longer small, a Taylor expansion in θ is exactly what is no longer available – the same loss of differentiability that makes mechanism (c) a genuine regime change rather than a large correction. The correct expectation is not an approximation around a central θ_0 at all, but the exact integral over the full nonlinear solution,

$$\mathbb{E}_{\theta\sim q}[y_x(\theta, z)] = \int g_x\big(\ell(\theta, z), z\big)\, q(\theta)\, d\theta, \qquad \ell(\theta, z) = F\big(\ell(\theta, z), z; \theta\big),$$

where $\ell(\theta, z)$ is not the linearised $(I - J(\theta))^{-1} B\delta z$ of Eq. (4) but its full, nonlinear fixed point, re-solved exactly at every θ in the support of the belief. There is no closed form: $g_x(\ell(\theta, z), z)$ is itself defined only implicitly, through a fixed point whose existence and uniqueness are exactly what mechanism (c) can take away. In the calm regime that fixed point is well-behaved and varies smoothly enough in θ that a point estimate $y_x(\theta_0, z)$ is a fair stand-in for $\mathbb{E}_{\theta\sim q}[y_x(\theta, z)]$; in the critical regime the only way to evaluate the expectation honestly is to resolve the network's full nonlinear fixed point at a representative sample of θ's drawn from q and average the results – the same Monte-Carlo-over-the-full-model prescription the risk calculation below already requires, now applied to uncertainty about the state itself rather than only about the shock.

Dynamics. Performance over a window of several periods is not a function of the current shock alone, and tracking it requires the exact, un-linearised loop, not a sensitivity carried forward in time. Unrolling the loop of Section 2.5,

$$Y_x = \sum_{t=0}^{T-1} y_x(\theta_t, z_t), \qquad y_x(\theta_t, z_t) = g_x(\ell_t, z_t), \qquad \ell_t = F(\ell_t, z_t; \theta_t), \qquad \theta_{t+1} = \theta(\ell_t, z_{t+1}, d_t),$$

each ℓ_t is the full nonlinear fixed point of Eq. (3) at period t's state, not a linear response carried over from $t-1$: no derivative of the loop is taken anywhere in this recursion, only its exact solution, re-computed period by period. Y_x is therefore a functional of the whole trajectory of actions $(d_0, \ldots, d_{T-1})$, not of the shock realised this period. A finance team attributing this quarter's EBIT miss entirely to this quarter's inflation, rate and equity moves is implicitly assuming $\theta_t \equiv \theta_0$ for every t in the window; once other firms' hoarding or pre-emptive ordering in period $t-1$ has already moved θ_t, part of this quarter's result is the delayed consequence of last quarter's actions, not of this quarter's shock, and no amount of re-running Superposition and Attribution on this period's z alone will recover it.

Strategy

The breakdown reaches strategic decisions in two separate ways, mirroring the breakdown of Δ_x itself. First, because Π_x depends on z through the same shock-dependent J as mechanism (b), the curvature terms behind Δ_x^d no longer stay put: $\mathrm{d}d_x^*/\mathrm{d}z$ measured at the baseline understates how much the optimal choice actually needs to move once a large shock arrives, exactly as a Δ_x measured in a calm state understates the sensitivity that will apply once the shock is large. Second, and more seriously, mechanism (c) means Π_x can lose differentiability altogether at a threshold, so the first-order condition $\partial_{d_x}\Pi_x = 0$ stops pinning down a unique, smoothly-moving d_x^*: the optimal choice can jump at a point where a plan built on Δ_x^d would have predicted nothing more than a gradual adjustment. A strategic plan set once against a central scenario has no way to anticipate this, because the jump is not a large value of a quantity the plan was tracking; it is the point at which the plan's own mathematical premise, differentiability, stops holding. In practice this shows up as a strategic plan revised by more than its own sensitivity analysis said it should, then abandoned outright rather than adjusted: a capacity-expansion plan calibrated on "if demand grows $X\%$, add $Y\%$ capacity" that turns out to need far more than Δ_x^d implied, followed by a sourcing strategy that

flips from single-sourcing to deliberate multi-sourcing almost overnight once a supplier disruption crosses a threshold, rather than drifting there gradually the way the smooth part of Δ_x^d would suggest.

Uncertainty. Under a non-degenerate belief, the optimal decision is not the one that maximises expected payoff at the mean state, and the gap is not a variance correction to be added on top of that payoff – it is a different functional of the belief altogether.

Dynamics. Rather than ask whether a single period's best responses settle into a fixed point, the question becomes a dynamic one. Whether the firm's policies (or set of actions) are optimized against the possible future paths the network can take. This becomes a question to explore strategies forward-looking, evaluating the firm's decisions in an uncertain future and consider the impact of the strategy.

Risk

(a) Amplification near criticality. With eigenvalues μ_q and right and left eigenvectors u_q, v_q of J,

$$\delta\ell = \sum_q u_q \, \frac{v_q^\top B \, \delta z}{1 - \mu_q}. \tag{11}$$

As the largest eigenvalue μ_1 approaches one, the term $q = 1$ dominates the sum: a single mode dominates. Every shock is projected onto the same direction u_1, which is concentrated on the club and on the firms that depend on it. Exposures to inflation, rates and equities, which were distinct in the calm state, become proportional to one common factor $v_1^\top B e_k$. Different shocks become strongly correlated: the effective dimension of risk collapses to one. To some extent, this simplifies the problem. It is no longer needed to track a multitude of variables. Finding the risk dimension that drives the full network may be sufficient. This is where techniques capturing this dimensional reduction will outshine complex language models.

(c) Thresholds. When a buffer is crossed, the response is no longer differentiable and no Taylor expansion, however many terms it keeps, exists across the threshold: Eq. (5) is simply not defined there. Losses then cascade, as in models of default contagion [22,23,24], and in coupled networks a small initial failure can propagate widely [25].

(d) Dependence between the shocks themselves. Shocks are not independent of the regime. The correlation between bond and equity returns depends on the inflation regime [26], and correlations between asset returns rise in downturns [27,28]. A regime-switching description [29], in which the joint distribution of δz depends on a hidden state, is more faithful than a single covariance matrix estimated once and used regardless of the state. In the stressed state both the mapping $\delta z \mapsto \delta y_x$ and the distribution of δz change together, and in the same direction: large, correlated shocks are more likely exactly when the network amplifies them most.

Value-at-risk and risk contributions. Mechanisms (a)–(d) are not only a statement about point forecasts; they overturn the risk-aggregation formulas of Eq. (9) in three separate ways. First, $\mathrm{VaR}_\alpha \approx z_\alpha \sigma(\delta y_x)$ was built on a fixed Δ_x, but mechanism (b) says Δ_x is no longer fixed: it grows with the shock itself, so a Δ_x measured today, in whatever state the network happens to be in today, understates the sensitivity that will apply once a large shock actually arrives. Second, $\sigma^2(\delta y_x) = \Delta_x^\top \Sigma_z \Delta_x$ was built on a fixed covariance matrix Σ_z, but mechanism (d) says Σ_z itself changes with the state, and in the state where it matters most for the tail, so a covariance matrix estimated in calm times is measured in the wrong regime for a stress calculation. Third, and most fundamentally, mechanism (c) makes $y_x(z)$ discontinuous in its derivatives at the point where a buffer is crossed: no Δ_x, computed however carefully, can represent a jump, so any VaR formula built from $\Delta_x^\top \Sigma_z \Delta_x$ places essentially no probability mass beyond that point, which is precisely where the tail loss is concentrated once the network switches branch (Figure 3b compares the two calculations). The risk contributions RC_k of Eq. (9) rely on the same linearity: they add up to the total only because a linear map is homogeneous of degree one in the shocks (Euler's theorem). Once thresholds and state-dependent coupling are present, $z \mapsto y_x$ is homogeneous of no degree at all, so contributions computed shock by shock no longer sum to the realised loss. Appendix D and Appendix E show numerically that the unexplained residual can dominate the loss. A meaningful VaR_α follows from drawing joint shocks from a distribution whose covariance and tail shape depend on the state, solving the full nonlinear network model at every draw, and reading off the quantile (as in Figure 3b for the illustrative network). One corollary is that hedges sized against individual risk factors start moving together when they were assumed independent, and a stress test built by shocking one factor at a time turns out to have no scenario resembling what actually happened.

Uncertainty. The calm regime's $\mathrm{VaR}_\alpha \approx z_\alpha \sigma(\delta y_x)$ integrates only over the shock z, at a fixed, assumed θ, and does so through the linear δy_x. Once θ itself is uncertain and correlated with z – a large shock makes a stressed state more likely, and a stressed state makes a given shock propagate further – the two cannot be integrated separately, and the mapping inside the integral can no longer be the linear δy_x either: the loss distribution a risk function should use is the joint one, built from the nonlinear performance $y_x(\theta, z) = g_x(\ell(\theta, z), z)$,

$$\mathbb{P}\big(y_x \le y\big) = \iint \mathbf{1}\{y_x(\theta, z) \le y\}\, p(z \mid \theta)\, q(\theta)\, dz\, d\theta, \qquad \mathrm{VaR}_\alpha = \inf\{y : \mathbb{P}(y_x \le y) \ge 1 - \alpha\}.$$

The distributions $p(z \mid \theta)$ and $q(\theta)$ are coupled under stochastic integration over the network states. Hence assuming a fully independent state θ misses dependence crucial for risk management. Careful integration of the probability distributions considering the dependence of the state on the shocks is required for proper risk management.

Dynamics. The calm regime treats Σ_z and ρ_{eff} as parameters estimated once and reused; the critical regime drives them through the network's own dynamics, and the right description of that drive is not a decay rate estimated by linearising about one point. Let

$$\Phi : \theta \longmapsto \theta\big(\ell(\theta, z),\ z,\ d^*(\theta)\big)$$

be the exact, un-linearised map from one period's state to the next. Once hoarding externalities are present, Φ is generically non-monotonic and can have more than one fixed point. The calm-regime state θ^*_{calm} and the stressed-regime state $\theta^*_{\text{stressed}}$ now read as two distinct fixed points of Φ. Regime switches are established when θ_t is swept out of the basin of attraction of θ^*_{calm} and into that of $\theta^*_{\text{stressed}}$. Once inside the stressed basin, a return of z to the calm regime does not guarantee that θ_t moves back into the calm basin, a phenomenon called hysteresis. Appendix E.8 documents it in Europe's 2022 energy crisis: gas prices fell by roughly 90% from their peak, yet energy-intensive production did not return to its former level, because the state had moved to a different fixed point of Φ, not merely to a different point on the same curve. A risk model that re-estimates Σ_z and ρ_{eff} from a rolling window of recent, already-stressed data is implicitly still assuming a single curve with one fixed point to track, which is exactly why it discovers only after the fact that it had quietly stopped working.

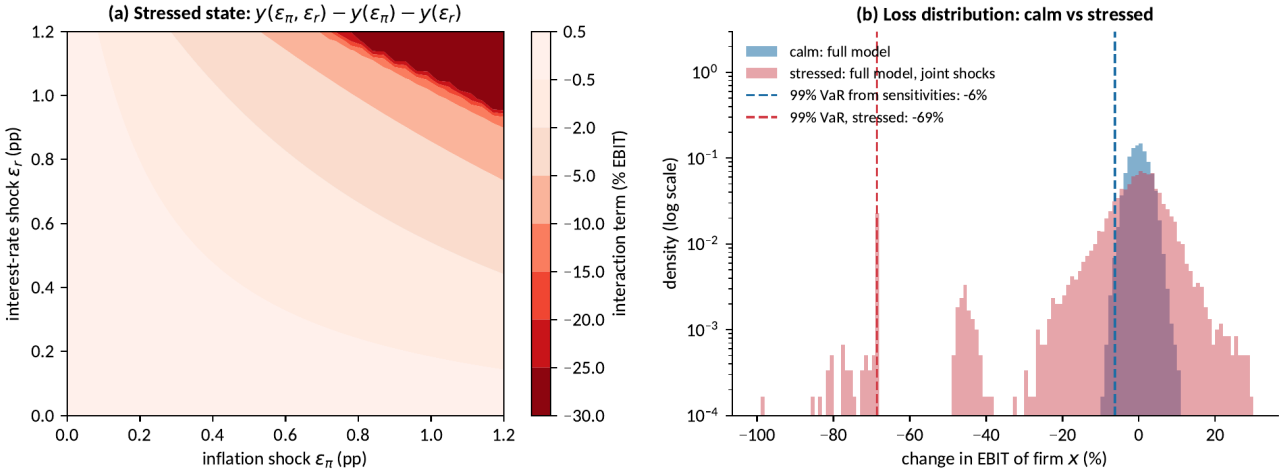


Figure 3. (a) Interaction term between inflation and interest-rate shocks for firm x in the stressed state; in the calm state it is close to zero everywhere. (b) Distribution of the change in EBIT of firm x under random joint shocks: calm state with moderate correlations, and stressed state with high correlations and fat tails (Student-t, 4 degrees of freedom). Dashed lines: 99% VaR from sensitivities (calm) and from the full model (stressed). Illustrative network of Appendix C.1.

In short, when $\kappa \ll 1$ no longer holds, strategy, financial performance and risk management cannot anymore ignore network effects because the response of x runs through $(I - J)^{-1}$ whose behaviour is set by the interconnected club. Proper management in this regime requires consideration of network dependencies, uncertainty of outcomes and coupled dynamics.

Appendix E contains a worked through numerical example. It reproduces the same regime switch with concrete magnitudes, and applies the same diagnostic to a real episode: the 2022 European energy and inflation crisis, in which the energy-intensive part of the economy crossed its critical point while the wider economy stayed below it, and never fully recovered.

3.3 What capturing criticality requires

Four requirements follow directly from the mechanisms above, and together they set what any framework for managing performance close to criticality has to be able to do.

First, it has to be non-perturbative. The classical toolkit of Section 3.1 is built on a linear expansion around a baseline: a small change is assumed to produce a proportionally small, predictable effect. Section 3.2 showed that this assumption fails exactly where it matters most – close to criticality the expansion stops converging, and no number of correction terms rescues it. A framework for this regime cannot treat the stressed state as a perturbation of the calm one; it has to reason directly about a state in which effects are no longer small, proportional or additive.

Second, it has to be a network model, not a single-firm model. The behaviours described above – hoarding that feeds on itself, a small action that ripples outward, adjustments that keep arriving faster than they can be resolved – are not properties of any one firm considered on its own; they are properties of how firms are coupled to one another through the structure set out in Sections 2 and 3. A framework that reasons about a single firm in isolation, however detailed, cannot capture conditions whose source is the network itself. It has to represent the network explicitly, and trace effects through it as they unfold, rather than compressing the network's influence into a single number computed once and reused.

Third, it has to capture joint, regime-dependent probability behaviour, rather than treating shocks and actions as independent and unconditional. Mechanism (d) of Section 3.2 showed that shocks which look independent in a calm state become correlated exactly when the network is under stress; the same is true of firms' actions here – whether a competitor hoards, raises a price or tightens a term is far from independent of what other firms are already doing once the network sits close to its critical point. A framework has to generate and reason about these possibilities jointly, conditioned on the state of the network, rather than drawing them one at a time from a fixed and unconditional distribution.

Fourth, the framework has to be dynamic. Joint shocks can swing a network from one regime into another. Proper management needs to capture these effects.

Active inference, introduced next, is built to meet all four requirements at once: it is non-perturbative by construction, it traces beliefs and actions through the network rather than through a single firm, it reasons about scenarios jointly rather than shock by shock, and it is dynamic, updating beliefs and actions period by period as the network moves.

4. **Scenario generation:** What generating scenarios means for strategy, performance and risk

The four requirements of the previous section have a direct consequence for how a firm should think about its future, before any particular framework is chosen to act on them.

1. **Strategy, opportunity, performance and risk become one exercise.** Once effects are non-perturbative, network-wide and jointly dependent on regime, "what happens to my risk", "what happens to my performance" and "where is the opportunity" stop being separate questions answered by separate teams with separate tools. They are read-outs of the same joint, network-propagated scenario. In Appendix E's calculated example, the scenario that threatens a 118% hit to the small supplier's profit is the same scenario in which a hub with spare capacity and intact buffers can win the business of weaker competitors; the loss side and the gain side come from one calculation through the network, not two independent ones.
2. **Scenario thinking has to replace single-driver thinking.** Because the network's response to several drivers moving together is not the sum of its responses to each moving alone (Section 3.2), a firm cannot assess its position by moving one rate, one price or one exchange rate at a time and adding up the results, however carefully each single-driver sensitivity is estimated. What has to be assessed is the joint, regime-dependent realisation of the relevant drivers.
3. **The set of scenarios considered has to be representative enough to avoid bias.** A network under stress generates a distribution of outcomes that is both skewed and regime-dependent. A small, hand-picked set of "plausible" scenarios tends to under-sample that correlated tail, producing an optimistic bias that grows, rather than shrinks, as the network approaches its critical point. In practice this means two things: scenarios have to be generated together, with a dependence structure that changes with the regime, rather than one driver at a time or from a single central forecast. The exercise should track where the network currently sits relative to criticality, since indicators such as the variance and autocorrelation of key performance indicators tend to rise as a critical point is approached [30]. That distance is what decides whether a simpler, sensitivity-based assessment can be trusted or not.
4. **This leads to a simple operating framework: generate scenarios, assess impact, and act.** Scenarios are generated jointly, across the relevant drivers, with a dependence structure conditioned on the network's current state rather than drawn one driver at a time. Their impact is assessed by propagating each scenario through the network – for strategy, performance and risk together. The firm acts on the resulting range of outcomes, choosing actions that are evaluated.

5. **Active inference:** A framework for measuring impact under network effects

Once a firm's problem has the shape described in Section 4, the right way to think about a firm's decision-making has more in common with how an organism copes with an uncertain environment than with a textbook optimisation exercise. Active inference is a framework, originating in neuroscience as a theory of how adaptive agents perceive and act under uncertainty, built around a loop of four steps that a firm goes through continuously.

1. Firms hold a belief, not a certainty, about their environment. A firm never knows for sure whether the network around it is calm or stressed, whether its suppliers have spare capacity, or how its competitors will behave if conditions tighten. What it has instead is a working view of the situation, built from everything it has seen so far. That belief can be written down as a causal picture of the firm and its environment. This is revised continuously as new evidence comes in, rather than fixed once at the start of a scenario.

2. Firms observe what other firms do. A firm cannot see the true state of the network directly, but it can see the traces other firms leave: order volumes, lead times, prices, payment behaviour, who else is chasing the same suppliers it is. These observable actions are its window onto a hidden situation, and it reads them as signals about whether the network is drifting towards a calm or a stressed state. The ability by which a firm is able to understand and adapt to the state of the network, influences the firm's competitive position.

3. Firms adapt their strategy. Given its current belief, the firm decides what to do – pricing, ordering, hedging, sourcing. A firm reasoning this way does not only chase the best outcome under its current belief. It also puts weight on actions that would tell it more about the future state of the network. Early information is worth paying for; Section 6 works out what this means concretely for hedges, performance targets and strategy.

4. Firms update their beliefs. As the consequences of its own actions and of others' actions play out, the firm compares what happened with what it expected, and revises its belief accordingly. The loop then repeats.

Why this is the right framework. What makes this useful for measuring impact under network effects is that active inference is, by construction, a *non-perturbative* framework: it does not assume that the effect of an external shock on a firm can be read off from a fixed sensitivity, and it does not assume that firms' reactions simply add up. Instead, it traces the actual chain of belief and action through the network – what a firm believes, what it does, how others respond, and what that implies in turn – to work out what a shock to the environment should mean for a firm's own strategy, exactly in the regime described in Section 3.2 where a fixed sensitivity vector stops being a reliable guide. It is precisely because active inference does not depend on the network staying in a stable, linear regime that it remains useful once perturbation theory has broken down.

At the same time, the framework does not discard the classical toolkit where that toolkit works. When the regime is known with confidence and beliefs are settled, there is nothing further to be gained from probing the environment or valuing information, and the firm's problem collapses back to simply choosing the action that is best given what it already knows – the quasi-equilibrium optimisation of Section 3.1, where sensitivities are sufficient. Active inference is best read as the general case, of which calm-state, sensitivity-based management is the special case that applies once the network is comfortably away from criticality.

Measuring impact this way is only useful if it changes what a firm does. Turning an assessed impact into a hedge, a performance target or a revised strategy is the subject of the next section.

6. **Acting:** Defining measures to hedge risk, optimise performance and refine strategy

Sections 3 to 5 show that a firm's problem changes nature, not only size, when the network moves from a calm to a stressed state, and give a way of seeing that change coming: generate scenarios jointly, and assess their impact through the network with active inference rather than fixed sensitivities. What is left is to act: to hedge risk, to optimise performance under uncertainty, and to define and refine strategy.

Hedging risk. Write the change in a unit's performance under a joint scenario as the sum of its sensitivity-based estimate and a network residual,

$$\delta y_i \;=\; \sum_k \Delta_{i,k}\,\varepsilon_k \;+\; \mathrm{Res}_i(\kappa),$$

where $\mathrm{Res}_i(\kappa)$ collects the interaction, second-order and regime-switching terms of Section 3.2 and Appendix B. For $\kappa \ll 1$, $\mathrm{Res}_i \approx 0$ and the ordinary sensitivity sum in the calm-state columns of Table 1 is both correct and sufficient to hedge against. As $\kappa \to 1$, Res_i is no longer small. Hedging is a matter of pricing $\mathrm{Res}_i(\kappa)$, not assuming it away. This has a simple economic logic: an option on substitutability, or on capacity, pays off mainly on the high-stress branch, where alternatives are scarce and contested, but it can only be bought cheaply on the low-stress branch. Its value is highest exactly when a sensitivity-based view says it is least needed, so hedges have to be put in place in calm periods, which is only possible if the firm can see the stressed branch before it is on it.

Optimising performance under uncertainty. Because a company's units are coupled through the network, the action that is best for one unit in isolation is not necessarily the action that is best once its effect on other units, through shared suppliers, customers and financing, is taken into account; Companies considering the actions of its units jointly in concurrence with the wider network, can form a holistic optimization of the set of actions that optimizes performance under uncertainty.

Defining and refining strategy. A small, deliberately informative order, a trial substitution, or a limited pilot with a new supplier can be worth placing purely because it shortens the time to a confident regime call, ahead of any scenario actually biting. This is what it means to define and refine strategy under this framework: not a single plan fixed in advance, but a sequence of actions holistically selected, each one feeding back into the belief of step 1, and each revision of that belief feeding forward into the next action. Strategy in an uncertain network becomes an adaptive and monitored set of actions pervading the company, its units and its environment.

7. **Introducing hedgeo:** Scenario intelligence for firms in networks

Hedgeo provides European scenario intelligence, built to help companies prepare for the regime described in this paper. It addresses sectors sensitive to network effects. Examples are logistics and shipping, industrial manufacturing and energy-intensive companies. It works through two branches: consulting, for a tailored analysis of risks and their interdependencies, and software, for the continuous and scalable version.

- **It defines the question to be answered in the company's own words.** A user states a risk, performance or strategy question. Hedgeo turns it into a set of structured, joint scenarios, expressed as values for macro variables such as oil prices, freight rates, the dollar, yield curves and other key variables to the company.
- **It represents the network as a causal graph.** The model is a directed acyclic graph over the company's key variables. Each unit of a company represents a subgraph.
- **It updates beliefs with active inference.** Beliefs about hidden drivers are updated as data arrive, and the platform reports the impact of a scenario on the company's key performance indicators against a baseline. The model is a readable document that can be inspected, not a black box.
- **It recommends action.** The platform proposes both financial hedging and business strategy for each downside scenario, with expected cost and benefit, which makes "buying substitutability in advance" concrete and comparable (Section 6).
- **It lets smaller firms benefit from sector knowledge without sharing their books.** The client's accounting stays within its own boundary, while sector-level structure helps a peripheral firm see the dependencies it cannot observe itself.
- **It is hosted in Europe**, in line with Europe's aim to build more resilience.

8. **Conclusion:** Managing under all circumstances requires scenario thinking, holistic impact measurement and acting under uncertainty.

Firms operate in networks of suppliers and customers, and in many economies those networks have a rich club of large firms at their core. For a firm in such a network, single-shock sensitivities are a correct and efficient basis for risk and performance management as long as the network is static and far from criticality.

That condition is not guaranteed. What counts as a small shock depends on the network's distance to criticality, and that distance shrinks when buffers are drawn down, substitutes are contested and credit is tight. In our calculated example (Appendix E), the same firms, network and scenario then produce losses of 118% and 42% of pre-tax profit instead of the 41% and 20% implied by sensitivities; shocks interact, the dependence between shocks changes, and competitors' dynamics become part of the problem. The 2022 energy crisis showed this in real data, where Europe's energy-intensive subgraph switched to a high-stress branch and stayed there, a persistence that does not show in perturbation theory but does under active inference. In this regime, firms need risk management that models dependencies between firms and between shocks, and a way of making decisions that infers the regime, values information and anticipates others. Active inference offers such a framework and contains the classical approach as its calm-state limit. Hedgeo is built on these principles. The future cannot be predicted, but a firm can see the stressed branch before it is on it, and prepare while preparation is still cheap.

Appendix: Notation, derivation and models

A. Notation

General network (Sections 2 and 3).

$\mathcal{H},\ \mathcal{P}$	rich club (hubs), periphery	$\phi(k)$	rich-club coefficient
$A,\ a_{ij}$	structural dependence of j on i	$s_{ij},\ a_{ij}^{\text{eff}}$	substitutable share; effective dependence
$\rho_{\text{eff}},\ \rho_0$	loop gain (spectral radius); at baseline	ℓ	activity losses of all firms
$z,\ \delta z$	macro variables and shock	$\xi(z),\ F$	direct exposure; propagation map, Eq. (3)
$J,\ B$	Jacobians of F in ℓ and z	$y_x,\ g_x$	performance indicator of firm x and its function
$\Delta_x,\ \Gamma_x$	sensitivities and second-order terms	$S_{\mathcal{H}}$	effective coupling of the club, Eq. (6)
$\kappa,\ c$	expansion parameter, Eq. (7); coupling sensitivity	$\Sigma_z,\ z_\alpha$	covariance of shocks; normal quantile
$\mu_q,\ u_q,\ v_q$	eigenvalues and eigenvectors of J	RC_k	risk contribution of shock k
$\varepsilon_k,\ e_k$	size and direction of shock k	Res_x	attribution (network) residual
$\theta,\ \Theta$	state of the network; set of possible states	$q(\theta),\ q_{i,t}$	a firm's belief about the state; firm i's belief at t
$d_i,\ \mathcal{D}_i$	decision of firm i; its set of possible decisions	d_{-i}	decisions of all other firms
$\Pi_i,\ \Delta_x^d$	objective of firm i; strategic sensitivity $\mathrm{d}d_x^*/\mathrm{d}z$	VaR_α	value-at-risk at confidence level α
Φ	one-period map of the network state	$\theta^*_{\text{calm}},\ \theta^*_{\text{stressed}}$	calm and stressed fixed points of Φ

Calculated example (Appendix E). Here the first-round loss ξ_i is the counterpart of the direct exposure $\xi(z)$ of Section 2, the shock vector ε that of the shock δz, and the coupling matrix G that of the Jacobian J.

$Z_{ij},\ R_i$	sales of i to j; revenue	$\gamma_{ij},\ c_{ij},\ \omega_i$	sales share, cost share, final-demand share, Eq. (15)
χ_{ij}	essentiality of supplier i for customer j	$s,\ s_0,\ s_f,\ u,\ u_c$	substitutability, its initial value and floor; system stress; stress where substitutes run out
$\varepsilon,\ \lambda$	shock vector; scenario intensity	$g_i,\ \xi_i$	final-demand response; first-round loss
$\hat{c}_j,\ \hat{p}_j,\ \beta_j$	unit-cost and price change; pass-through	$G,\ T$	coupling matrix; loss map, Eq. (17)
$\rho(G),\ \rho_T$	loop gain; local loop gain	N, E, O, W, K	network inputs, energy, other materials, variable wages, fixed costs
$D,\ f,\ r$	debt, floating share, interest rate	$\mathrm{EBIT},\ \mathrm{EBT},\ \mathrm{ICR}$	operating profit, pre-tax profit, interest cover
$sp,\ \eta$	credit spread; its sensitivity	$F_S,\ v,\ m_c,\ \sigma_m$	SME failures, vulnerable share, critical margin, dispersion
Liq_x	liquidity buffer of firm x	$\Delta_{i,k},\ \Gamma_{i,km}$	sensitivities; second-order terms

B. The expansion parameter

Let the Jacobian depend on the shock, $J = J_0 + J'\delta z$, and write $M_0 = (I - J_0)^{-1}$. Then

$$(I - J_0 - J'\delta z)^{-1} = M_0 \sum_{n=0}^{\infty} \left(J'\delta z\, M_0\right)^n,$$

which converges when $\|J'\delta z\|\ \|M_0\| < 1$. Near criticality $\|M_0\|$ scales as $1/(1-\rho_0)$, which gives Eq. (7) with c of the order of $\|J'\|$. A scalar version makes the structure transparent. Let $y = (\bar{\xi}_1\varepsilon_1 + \bar{\xi}_2\varepsilon_2)/(1 - \rho_0 - c_1\varepsilon_1 - c_2\varepsilon_2)$, where the $\bar{\xi}_k$ are direct exposures to shock k and the c_k measure how much each shock erodes substitutability. With $\kappa = (c_1\varepsilon_1 + c_2\varepsilon_2)/(1-\rho_0)$,

$$y = \frac{\bar{\xi}_1\varepsilon_1 + \bar{\xi}_2\varepsilon_2}{1-\rho_0} \cdot \frac{1}{1-\kappa} = \frac{\bar{\xi}_1\varepsilon_1 + \bar{\xi}_2\varepsilon_2}{1-\rho_0}\left(1 + \kappa + \kappa^2 + \dots\right).$$

In the calculated example, J is the coupling matrix G of Eq. (17). The first-order (sensitivity) result is the first term; the n-th correction is smaller by κ^n. The series diverges at $\kappa = 1$, and the second-order cross-term gives Eq. (10).

C.1 The illustrative network (Sections 2 and 3)

The network has 12 large and 150 small firms. Links among large firms exist with probability 0.8 (weights uniform on $[0.5, 1.5]$). Each small firm links to one, two or three large firms with probabilities 0.45, 0.40 and 0.15; its dependence on each has weight uniform on $[1, 3]$, and the large firm's dependence on it has weight uniform on $[0.05, 0.3]$. Links among small firms exist with probability 0.012 (weights uniform on $[0.3, 1]$). Columns are scaled so that total dependence is 0.85 for large and 0.90 for small firms. Direct loadings on (π, r, e) are uniform on $[0.004, 0.010]$ for large and $[0.004, 0.014]$ for small firms; firm x has $(0.0065, 0.005, 0.002)$. The substitutable share is $s(u, \varepsilon) = s_0/(1 + e^{(u-u_0)/w}) - c_r \max(\varepsilon_r, 0)$, where u is the size-weighted average activity loss (large firms 20 times the size of small ones). A firm whose loss exceeds its threshold τ loses at least 30% of activity. The change in EBIT is $-(1 - e^{-4\,\delta\ell_x})$, an operating leverage of four. Calm state: $s_0 = 0.70$, $u_0 = 0.20$, $w = 0.04$, $c_r = 0$, $\tau = (0.60, 0.45)$ for (large, small). Stressed state: $s_0 = 0.60$, $u_0 = 0.05$, $w = 0.05$, $c_r = 0.08$, $\tau = (0.30, 0.20)$, baseline stress of 0.5% activity for all firms. Equation (3) is solved by damped fixed-point iteration. For Figure 3b, 6,000 joint shocks are drawn: in the calm state Gaussian with standard deviation 0.45 units and correlations 0.3 (inflation, rates), 0.0 (inflation, equities) and 0.1 (rates, equities); in the stressed state Student-t with 4 degrees of freedom, scale 0.55 and correlations 0.7, 0.5 and 0.6. Results are mainly for the purpose of illustrating the argument.

C.2 The calculated example (Appendix E)

Data are in Tables 2–4. The essentialities χ_{ij} (share of customer j's output that needs supplier i) are: for AutoCo, ChemCo 0.25, SME periphery 0.35, LogiCo 0.15; for ChemCo, LogiCo 0.20, SME periphery 0.10; for LogiCo, AutoCo 0.05, ChemCo 0.05, SME periphery 0.15; for the SME periphery, ChemCo 0.25, LogiCo 0.15, AutoCo 0.05; for firm x, ChemCo 0.15, SME periphery 0.25, LogiCo 0.15. Prices follow Eq. (16); the loss map (17) depends on substitutability s, which erodes with system stress u (the revenue-weighted average volume loss),

$$s(u) = s_0\ \max\left(s_f,\ 1 - \frac{u}{u_c}\right), \qquad u = \frac{\sum_i R_i \ell_i}{\sum_i R_i}, \tag{12}$$

where s_0 is the substitutability in the undisturbed state, u_c the stress level at which substitutes run out, and s_f the fraction of

s_0 that always remains. Two threshold mechanisms complete the model. A share v of the SME periphery has thin buffers and fails once its pre-tax margin m_S falls below a critical level m_c; with margins dispersed around m_S with standard deviation σ_m, the additional failure share caused by the shock is

$$F_S = v\Big[\mathcal{N}\Big(\frac{m_c - m_S'}{\sigma_m}\Big) - \mathcal{N}\Big(\frac{m_c - m_S}{\sigma_m}\Big)\Big], \tag{13}$$

with $\mathcal{N}$ the standard normal distribution function and m_S' the margin after the shock; failed firms that are not replaced (a share $1-s$) add to the loss of the node, $\ell_S \to 1-(1-\ell_S)(1-F_S(1-s))$, and firm x breaches its liquidity buffer Liq_x once its pre-tax loss exceeds it. Separately, when a firm's interest cover $\mathrm{ICR} = \mathrm{EBIT}/I$ falls, its credit spread widens by

$$\Delta sp_i = sp_{i,0}\Big[\Big(\frac{\mathrm{ICR}_{i,0}}{\mathrm{ICR}_i'}\Big)^{\eta} - 1\Big], \tag{14}$$

with $sp_{i,0} = r_i - 2\%$ the initial spread over a 2% base rate and $\eta \geq 0$ the sensitivity of lenders to interest cover ($\eta = 0$: spreads do not react). The loss map (17) with Eqs. (12)–(14) is solved by repeated substitution, starting from $\ell = 0$ and computing in each pass system stress, substitutability, the coupling matrix, the new losses, the SME periphery's pre-tax margin and its failure share, for 300 passes. Pre-tax profit then follows from Eq. (18). Sensitivities are numerical derivatives of the full model at $\varepsilon = 0$ with step 10^{-5}; the direct part holds other firms' prices and volumes fixed. The local loop gain is the spectral radius of the numerical Jacobian of one pass of the loss map at the solution.

D. A worked numerical illustration of Section 3

This appendix puts numbers on Section 3, on the illustrative network of Appendix C.1. Take firm x to be a small industrial supplier in the periphery that sells to two large firms, say a manufacturer and a retailer, and buys energy and materials. One unit of the scenario is $+1$ percentage point (pp) of inflation, $+1$ pp of interest rates and a 10% fall in equity markets. Each shock reaches x through a direct channel and through the network:

- **Inflation.** Directly, input costs rise faster than contracted selling prices. Through the network, suppliers pass on their own cost increases, while the large customers use their bargaining power to delay price increases, which squeezes the margin of x from both sides.
- **Interest rates.** Directly, floating-rate debt and refinancing become more expensive. Through the network, large customers cut investment and orders, and lengthen payment terms, which raises the working capital that x must finance.
- **Equity markets.** The direct effect is small (pension funding, cost of equity). Through the network, lower household wealth reduces consumer demand at the retailer, and this travels upstream to its suppliers.

Calm state. The calm-state columns of Table 1 show the result. The loop gain is $\rho_{\text{eff}} = 0.26$. Between 30% and 57% of each sensitivity, and about 40% of the total, comes through the network, which confirms that a firm-level sensitivity is already a network quantity. The sum of the three sensitivities, -8.8% of EBIT, is within 0.3 pp of the full model, -8.5%, and the difference comes from the curvature of EBIT in activity (operating leverage), not from the network. The interaction between inflation and rates is $+0.1$ pp. For risk aggregation, with calm-state volatilities and correlations (Appendix C.1), the 99% VaR from Eq. (9) is -6.3% of EBIT against -6.0% in the full model (Figure 3b): in this state, sensitivity-based VaR and risk contributions are reliable, and firm x can manage the scenario with its sensitivities.

	Calm state ($\rho_{\text{eff}} = 0.26$)			Stressed state ($\rho_{\text{eff}} = 0.51$)	
Change in EBIT of firm x (%)	direct	via network	sensitivity	sensitivity	full, one shock
Inflation $+1$ pp	-2.6	-1.1	-3.7	-6.1	-6.6
Interest rates $+1$ pp	-2.0	-1.2	-3.2	-6.3	-8.5
Equity markets -10%	-0.8	-1.1	-1.9	-4.2	-4.8
Sum of single shocks	-5.4	-3.4	-8.8	-16.6	-19.9
Full model, joint scenario			-8.5		-68.6 (buffer breached)
Residual (joint minus sum)			$+0.3$	-52.0	-48.7
Interaction inflation × rates			$+0.1$	-9.1	

Table 1. Firm x under the scenario of $+1$ pp inflation, $+1$ pp interest rates and -10% equities, in two states of the same network. "Sensitivity" is the first-order response, Eq. (5); "full" solves the nonlinear model, Eq. (3). Illustrative network of Appendix C.1.

Stressed state. Now keep the firm, the network and the scenario, and change only the state: buffers have been drawn down by earlier shocks, some of the alternative suppliers are already taken by other firms (s falls from 0.70 to 0.60), substitutability erodes faster as system stress rises, and credit is tight, so a rise in rates makes switching supplier harder. The loop gain doubles to $\rho_{\text{eff}} = 0.51$ (Table 1, stressed columns, and Figure 2b). The sensitivities measured in the stressed state are already 1.6 to 2.2 times as large as in the calm state, which is the first sign of trouble. More importantly, they no longer describe the outcome. For small intensities the full model follows the expansion, but it bends away quickly. At $\lambda \approx 0.77$ it jumps: substitutability collapses and the network moves from the low-stress to the high-stress branch, the hysteresis described in [1]. At $\lambda \approx 0.85$ the periphery breaches its buffers. At the full scenario, $\lambda = 1$, EBIT falls by 68.6%, against 16.6% from the stressed sensitivities and 19.9% from the sum of the single-shock responses. The large firms do not fail; the club transmits the stress, and the periphery absorbs it.

By mechanism (b) of Section 3.2, the interaction between inflation and rates (Figure 3a) is -9.1 pp of EBIT in the stressed state against $+0.1$ pp in the calm state – a hundredfold increase for the same pair of shocks. By mechanism (d), with correlated, fat-tailed joint shocks in the stressed state (Appendix C.1), the 99% loss of firm x is about 69% of EBIT (Figure 3b), against 6% from the sensitivity-based calculation of Eq. (9): the sensitivity-based VaR misses almost the entire tail. In short, all four properties of Section 3.1 fail here: superposition fails because of the interaction just computed, attribution fails because the residual (-52.0 pp out of -68.6 pp) dominates the sum of single-shock contributions, aggregation and risk contributions fail because the calm-state covariance and sensitivities used in Eq. (9) no longer describe the stressed state, and a hedge sized on the calm-state Δ_x would have covered only a small fraction of the eventual loss.

E. The calculated example: a small supplier, a large hub and a real crisis

Sections 2 and 3, and Appendix D above, made the argument on an illustrative network. This appendix builds a small network in which every number can instead be traced to an income statement, so the same argument can be checked by hand, in a spreadsheet or in code, and applies the same diagnostic to a real episode, the 2022 European energy and inflation crisis. We keep the rich-club structure but aggregate most of the periphery into one node.

E.1 From dependence shares to money flows

As in Section 2, firms i and j belong to a rich club $\mathcal{H}$ of large firms (the hubs) or to a periphery $\mathcal{P}$ of small firms. Here links are measured in money. We write Z_{ij} for the value of the goods and services that firm i sells to firm j in a year, and R_i for the revenue of firm i. Two ratios describe how much two firms depend on each other, one from each side of the link:

$$\gamma_{ij} = \frac{Z_{ij}}{R_i}, \qquad c_{ij} = \frac{Z_{ij}}{R_j}, \qquad \omega_i = \frac{R_i - \sum_j Z_{ij}}{R_i}. \tag{15}$$

The *sales share* γ_{ij} is the share of supplier i's revenue that comes from customer j; it measures how much i depends on j as a customer. The *cost share* c_{ij} is the share of customer j's revenue spent on inputs from i. The *final-demand share* ω_i is the share of i's revenue sold to households, governments and foreign buyers outside the network. Cost shares alone can understate how much an input matters, so we also record, for each supply link, an *essentiality* $\chi_{ij} \in [0, 1]$ (the share of customer j's output that stops if supplier i cannot deliver and no alternative is found) and a *substitutability* $s \in [0, 1]$ (the share of such a shortfall that other suppliers, inventory or spare capacity can replace); Appendix E.3 lets s fall as the network gets more stressed.

In the notation of Section 2, the effective dependence of firm j on firm i is now built from money flows: $a^{\text{eff}}_{ij} = \gamma_{ji} + (1 - s)\chi_{ij}$, combining the demand channel (firm j sells to i) and the supply channel (firm j needs i's input). The coupling matrix used below is its transpose, $G = (A^{\text{eff}})^{\top}$, and the macro shock δz of Section 2 becomes ε.

E.2 The example network

To make the argument concrete, we build a network of five nodes with realistic income statements (Tables 2 and 3, Figure 4). The club consists of three large firms: *AutoCo*, a car manufacturer with revenue of €150 bn; *ChemCo*, a chemicals producer with €60 bn; and *LogiCo*, a logistics group with €25 bn. They trade heavily with each other: ChemCo sells €6 bn of coatings and plastics to AutoCo, LogiCo sells €3 bn of transport to AutoCo and €2.5 bn to ChemCo, and AutoCo sells €1.5 bn of vans and trucks to LogiCo. Five of the six possible directed links among the three hubs exist, so the club has a link density of $5/6$. The periphery has two nodes. *Firm* x is a small precision-parts supplier with revenue of €60 m, of which 55% is sold to AutoCo and 25% to LogiCo. The *SME periphery* node represents the many other small and medium-sized suppliers of the club, with combined revenue of €40 bn, of which €30 bn is sold to AutoCo. Treating them as one node keeps the example small while preserving the fact that the club depends on the periphery collectively.

Seller i \ buyer j	AutoCo	ChemCo	LogiCo	SME per.	Firm x	Final demand	Revenue R_i
AutoCo (A)	–	0	1,500	1,000	0	147,500	150,000
ChemCo (C)	6,000	–	300	8,000	3	45,697	60,000
LogiCo (L)	3,000	2,500	–	3,000	1.5	16,498.5	25,000
SME periphery (S)	30,000	4,000	3,000	–	6	2,994	40,000
Firm x (X)	33	0	15	6	–	6	60

Table 2. Flows Z_{ij} between the five firms (EUR million per year): row i sells to column j. Final demand is revenue minus sales to the network.

	AutoCo	ChemCo	LogiCo	SME per.	Firm x
Revenue R	150,000	60,000	25,000	40,000	60.0
Network inputs N	39,033	6,500	4,815	12,006	10.5
Energy E	4,500	9,000	5,000	1,600	3.0
Other external materials O	61,500	21,000	4,000	10,400	17.0
Variable wages W	12,000	5,000	6,500	8,000	12.0
Fixed costs K	21,000	13,500	3,200	5,600	12.3
EBIT (operating profit)	11,967	5,000	1,485	2,394	5.20
Interest $I = rD$	1,000	700	320	800	1.32
Pre-tax profit (EBT)	10,967	4,300	1,165	1,594	3.88
Debt D; floating share f	25,000; 30%	20,000; 40%	8,000; 50%	16,000; 60%	24; 50%
Interest rate r	4.0%	3.5%	4.0%	5.0%	5.5%
Pass-through β	0.6	0.7	0.8	0.5	0.3
Final-demand share ω	0.983	0.762	0.660	0.075	0.100
Final demand: inflation, rates, equities (%)	−0.5 / −2.0 / −0.6	−0.2 / −0.8 / −0.2	−0.2 / −0.5 / −0.2	−0.2 / −1.0 / −0.3	−0.2 / −1.0 / −0.3

Table 3. Income statements and behavioural parameters (EUR million unless stated). Network inputs are the column sums of Table 2. Final-demand responses are per scenario unit.

The income statements follow typical cost structures: an operating margin of 8% for AutoCo, 8.3% for ChemCo, 6% for LogiCo and the SME periphery and 8.7% for firm x. Costs are split into network inputs $N_i = \sum_k Z_{ki}$, external energy E_i, other external materials O_i, variable wages W_i, which move with volume, and fixed costs K_i (fixed staff, depreciation, overheads), which do not. Earnings before interest and taxes are $\mathrm{EBIT}_i = R_i - N_i - E_i - O_i - W_i - K_i$, interest is $I_i = r_i D_i$ with debt D_i and average rate r_i, and pre-tax profit is $\mathrm{EBT}_i = \mathrm{EBIT}_i - I_i$, the performance indicator used throughout. The pass-through β_i is the share of a unit-cost increase that firm i can pass on to its customers; the small firm x, which sells most of its output to AutoCo, has the weakest bargaining position ($\beta = 0.3$). The essentialities χ_{ij} are listed in Appendix C.2: for example, 35% of AutoCo's output needs parts from the SME periphery, 25% needs ChemCo's materials and 15% needs LogiCo's transport.

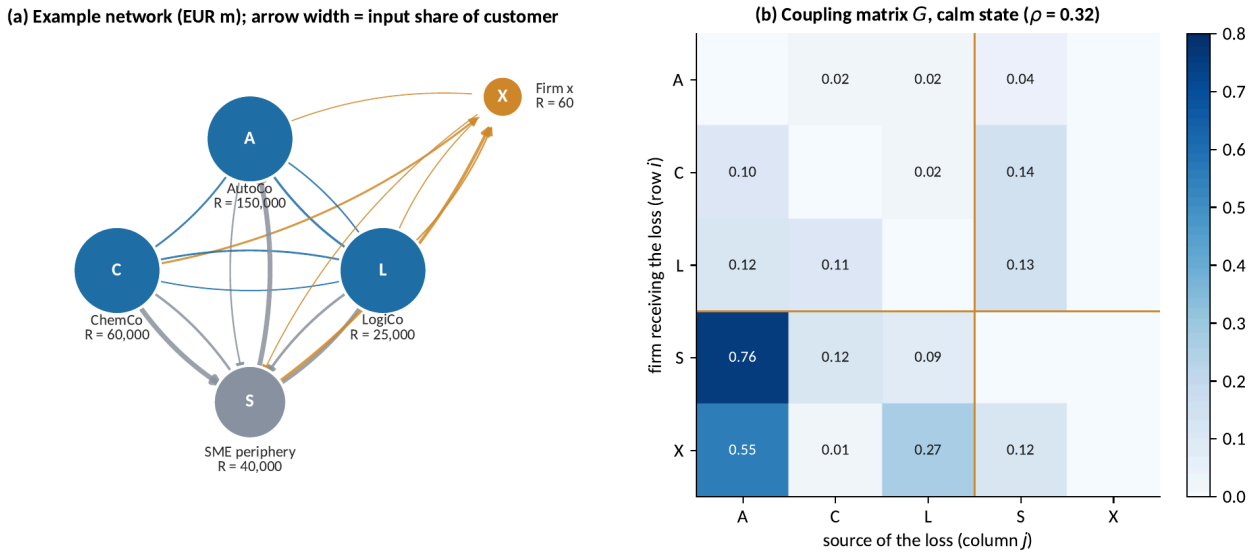


Figure 4. (a) The example network of Table 2; arrow width is proportional to the cost share c_{ij} of the input for the customer. (b) Coupling matrix G of Eq. (17) in the calm state: entry (i, j) is the fraction of a volume loss at firm j that is passed to firm i. The club block (top left) is dense; firm x receives 0.55 of AutoCo's losses.

E.3 How shocks travel through the network

Shocks. The macro-financial environment is summarised by a shock vector $\varepsilon = (\varepsilon_\pi, \varepsilon_r, \varepsilon_e)$ for inflation, interest rates and equity markets. One *scenario unit* is $+1$ percentage point (pp) of inflation, $+1$ pp of interest rates and a 10% fall in equities, and the *scenario intensity* λ scales it: $\varepsilon = \lambda(1, 1, 1)$. Each unit of inflation raises energy prices by 5%, other external materials by 1%, wages by 0.5% and fixed costs by 0.5% within the year. The large energy multiple reflects that energy prices move several times faster than the headline index; in 2022 euro area energy inflation was a multiple of headline inflation [31]. Each unit of interest rates adds 1 pp to the cost of floating-rate debt. Each shock also reduces final demand for firm i's products by $g_i = g_i^\pi \varepsilon_\pi + g_i^r \varepsilon_r + g_i^e \varepsilon_e$ (last row of Table 3): inflation through lower real income, interest rates through credit-financed purchases, which is why AutoCo is the most rate-sensitive at -2% per pp, and equities through household wealth.

Prices. Cost increases travel downstream. The change in unit variable cost of firm j, expressed as a share of its price, is its suppliers' price changes weighted by cost shares plus its own external cost increases, and it passes a share β_j on:

$$\hat{c}_j = \sum_i c_{ij}\, \hat{p}_i + \frac{0.05E_j + 0.01O_j + 0.005W_j}{R_j}\, \varepsilon_\pi, \qquad \hat{p}_j = \beta_j\, \hat{c}_j. \tag{16}$$

Here $\hat{p}_j$ is the relative change in firm j's selling price. The unpassed part, $(1 - \beta_j)\hat{c}_j$, comes out of its margin.

Volumes. Volume losses travel both ways. Let $\ell_i \geq 0$ be the relative fall in firm i's sales volume. It has three sources: the first-round loss of final demand $\xi_i = -\omega_i g_i$; the *demand channel*, in which a customer j that sells less buys less, so that i loses $\gamma_{ij}\ell_j$; and the *supply channel*, in which a supplier i that cannot deliver stops the essential share of its customer's output that cannot be substituted, $(1 - s)\chi_{ij}\ell_i$. Collecting terms gives the loss map

$$\ell = T(\ell) := \xi + G(s)\,\ell, \qquad G_{ij}(s) = \gamma_{ij} + (1 - s)\,\chi_{ji}. \tag{17}$$

The *coupling matrix* G says which fraction of a volume loss at firm j reaches firm i (Figure 4b). Its spectral radius $\rho(G)$, the largest absolute eigenvalue, is the network's *loop gain*: the factor by which a disturbance is multiplied each time it goes around the network. Substitutability erodes with stress as in Eq. (12) (Appendix C.2).

Buffers and credit. Two smaller mechanisms complete the model, mainly so the stressed state can show a genuine threshold rather than only a smooth erosion of substitutability. A share of the SME periphery has thin buffers and fails once its margin falls below a critical level; failed firms that are not replaced add to the node's volume loss. Firm x has a liquidity buffer and breaches it once its pre-tax loss exceeds it. Separately, when a firm's interest cover falls, lenders widen its credit spread, which adds a further interaction between the inflation and interest-rate shocks. The formulas for both are in Appendix C.2, Eqs. (13)–(14); none of the three add much to the argument on its own, and what matters for Section 3.2 is only that thresholds of this kind exist somewhere in a real network.

Performance. Given prices $\hat{p}$ and losses ℓ, pre-tax profit after the shock is

$$\begin{aligned}\mathrm{EBT}_i' = {} & R_i(1+\hat{p}_i)(1-\ell_i) - \sum_k Z_{ki}(1+\hat{p}_k)(1-\ell_i) \\ & - \big[E_i(1+0.05\varepsilon_\pi) + O_i(1+0.01\varepsilon_\pi) + W_i(1+0.005\varepsilon_\pi)\big](1-\ell_i) \\ & - K_i(1+0.005\varepsilon_\pi) - \big[r_i D_i + f_i D_i\,\Delta r + D_i\,\Delta sp_i\big].\end{aligned} \tag{18}$$

The first two lines are revenue minus variable costs; the third subtracts fixed costs and interest, with f_i the floating-rate share of debt and $\Delta r = \varepsilon_r$ pp.

E.4 Linear response, sensitivities and the role of the network

If s is held fixed and no buffer is crossed, the loss map (17) is linear and its solution is a series:

$$\ell = (I - G)^{-1}\xi = \sum_{k=0}^{\infty} G^k \xi, \qquad \rho(G) < 1. \tag{19}$$

The k-th term is the loss arriving after k steps through the network: directly ($k=0$), from customers and suppliers ($k=1$), from their customers and suppliers ($k=2$), and so on. A loss that goes around the network repeatedly is multiplied by at most $1/(1-\rho(G))$. For a firm i, expanding Eq. (18) in the shocks gives

$$\Delta\mathrm{EBT}_i = \sum_k \Delta_{i,k}\,\varepsilon_k + \tfrac{1}{2}\sum_{k,m}\Gamma_{i,km}\,\varepsilon_k\varepsilon_m + O(\|\varepsilon\|^3), \tag{20}$$

where $\Delta_{i,k} = \partial\mathrm{EBT}_i/\partial\varepsilon_k$ is the *sensitivity* of firm i to shock k and $\Gamma_{i,km}$ are second-order terms. We split each sensitivity into a *direct* part, computed with all other firms' prices and volumes held fixed, and a *network* part, the remainder, which arrives through Eqs. (16) and (19). The network part is why a firm's sensitivity is a network quantity.

E.5 Measuring the distance to criticality

The validity condition of Section 2.6 applies unchanged: the expansion (20) is reliable when the expansion parameter κ of Eq. (7) is small, with $\rho_0 = \rho(G)$ before the shock. In this example the most useful single number is the *local loop gain* $\rho_T := \rho(\partial T/\partial\ell)$, the spectral radius of the Jacobian of the full loss map (17) at the current solution. It includes the effect of eroding substitutability and of buffers, it equals $\rho(G)$ when neither is active, and it reaches one exactly where the solution stops depending smoothly on the shock, that is where κ reaches one.

E.6 Calm state: inflation up, interest rates up, equity markets down

Section 3.1 showed that in a calm regime superposition, attribution, aggregation and hedging all work with single-shock sensitivities. This subsection checks that on the calculated example for both firms; E.7 then changes the state of the network, and E.8 applies the same yardstick to a real crisis, including a numerical simulation of it. Table 4 lists the parameters of the two states of the network. In the calm state, 90% of an essential input shortfall can be substituted ($s_0 = 0.9$), substitutes only run out at 50% system stress, lenders do not react to interest cover ($\eta = 0$), firm x holds €8 m of cash and credit lines, and vulnerable SMEs fail only below a -4% margin. The loop gain is $\rho(G) = 0.32$, so a disturbance that circulates is amplified by at most $1/(1-\rho(G)) \approx 1.46$.

Symbol	Meaning	Calm	Stressed
s_0	substitutable share of an essential input shortfall	0.90	0.50
u_c	system stress at which substitutes run out	50%	10%
s_f	minimum fraction of s_0 that remains	0.2	0.2
η	sensitivity of credit spreads to interest cover	0	1
Liq_x	liquidity buffer of firm x (EUR m)	8.0	0.5
v	share of SMEs with thin buffers	15%	15%
m_c	pre-tax margin at which a vulnerable SME fails	-4%	-2%
σ_m	dispersion of SME margins	3%	3%
$\rho(G)$	loop gain before the shock	0.32	0.59

Table 4. State parameters. The firms, the network and the scenario are identical in both states; only buffers, substitutability and credit conditions differ.

Each shock reaches the two firms through direct and network channels, which Table 5 quantifies. Two calculations show how the numbers arise.

- **AutoCo, inflation, direct.** One unit of inflation raises AutoCo's energy bill by $0.05 \times 4{,}500 = 225$, other materials by $0.01 \times 61{,}500 = 615$, variable wages by 60 and fixed costs by 105: €1,005 m in total. AutoCo passes on $\beta = 0.6$ of its variable unit-cost increase, $0.6 \times (225 + 615 + 60)/150{,}000 = 0.36\%$ of its price, which brings in €540 m. Final demand falls by $0.983 \times 0.5\% = 0.49\%$, which costs 0.49% of the contribution margin of €32,967 m, that is €162 m. The direct sensitivity is $-1{,}005 + 540 - 162 = -627$.
- **Firm x, interest rates, network.** One unit of rates reduces AutoCo's final demand by 2%, which becomes a volume loss of about 2% at AutoCo after network feedback. Firm x sells 55% of its output to AutoCo, so it loses about 1.1% of volume on that account, plus smaller losses via LogiCo and the SME periphery. On a contribution margin of €17.5 m this costs €0.29 m, twice the direct effect of €0.14 m (of which €0.12 m is the extra interest on €12 m of floating-rate debt).

	Firm x (EBT 3.88)			AutoCo (EBT 10,967)		
	direct	network	sensitivity	direct	network	sensitivity
Inflation $+1$ pp	-0.33	-0.12	-0.45	-627	-94	-721
Interest rates $+1$ pp	-0.14	-0.29	-0.43	-723	-43	-766
Equities -10%	-0.01	-0.09	-0.09	-195	-13	-207
Sum of sensitivities	-0.47	-0.50	-0.98 (-25.1%)	$-1{,}545$	-150	$-1{,}695$ (-15.5%)
Full model, joint scenario			-0.99 (-25.6%)			$-1{,}713$ (-15.6%)
Residual; of which second order			-0.02; -0.02			-18; -17
Interaction inflation $\times$ rates			0.00			$+2$

Table 5. Calm state: change in pre-tax profit for one scenario unit (EUR million; in brackets % of base pre-tax profit).

The two firms differ in where their risk comes from. For firm x, 51% of its total sensitivity arrives through the network, mainly through AutoCo: the small firm's profit depends on the volume of the hub it supplies. For AutoCo, only 9% comes through the network; its exposure is dominated by its own final demand. The sum of the three sensitivities is within 0.5 percentage points of the full model for firm x (-25.1% against -25.6%) and within 0.2 percentage points for AutoCo (-15.5% against -15.6%). The residual is almost entirely the second-order term of Eq. (20), which comes from the product of price and volume changes, and the interaction between inflation and rates is negligible. Figures 5a and 5c show that this remains true up to twice the scenario. In the calm state both firms can manage the scenario with their sensitivities.

E.7 Stressed state: the same firms, network and shock

Now keep the firms, the network and the scenario, and change only the state (Table 4). This is the situation after a sequence of earlier shocks: alternative suppliers are already booked by other firms ($s_0 = 0.5$) and run out quickly as stress rises ($u_c = 10\%$); lenders reprice debt when interest cover falls ($\eta = 1$); firm x has used most of its liquidity ($\mathrm{Liq}_x =$ €0.5 m); and vulnerable SMEs fail already below a -2% margin. The loop gain before the shock rises from 0.32 to 0.59, and the maximum amplification from 1.46 to 2.41.

	Firm x		AutoCo	
	sensitivity	single shock	sensitivity	single shock
Inflation $+1$ pp	-0.60	-0.62	-825	-839
Interest rates $+1$ pp	-0.80	-1.03	$-1{,}090$	$-1{,}310$
Equities -10%	-0.18	-0.19	-291	-303
Sum	-1.58 (-40.6%)	-1.84 (-47.4%)	$-2{,}207$ (-20.1%)	$-2{,}452$ (-22.4%)
Full model, joint scenario	-4.58 (-118%); buffer breached		$-4{,}580$ (-41.8%)	
Residual (joint minus sum of sensitivities)	-3.01		$-2{,}373$	
Interaction inflation $\times$ rates	-0.42		-359	
Share of sensitivity via network	61%		26%	

Table 6. Stressed state: change in pre-tax profit for one scenario unit (EUR million; in brackets % of base pre-tax profit). "Single shock" solves the full model for each shock alone.

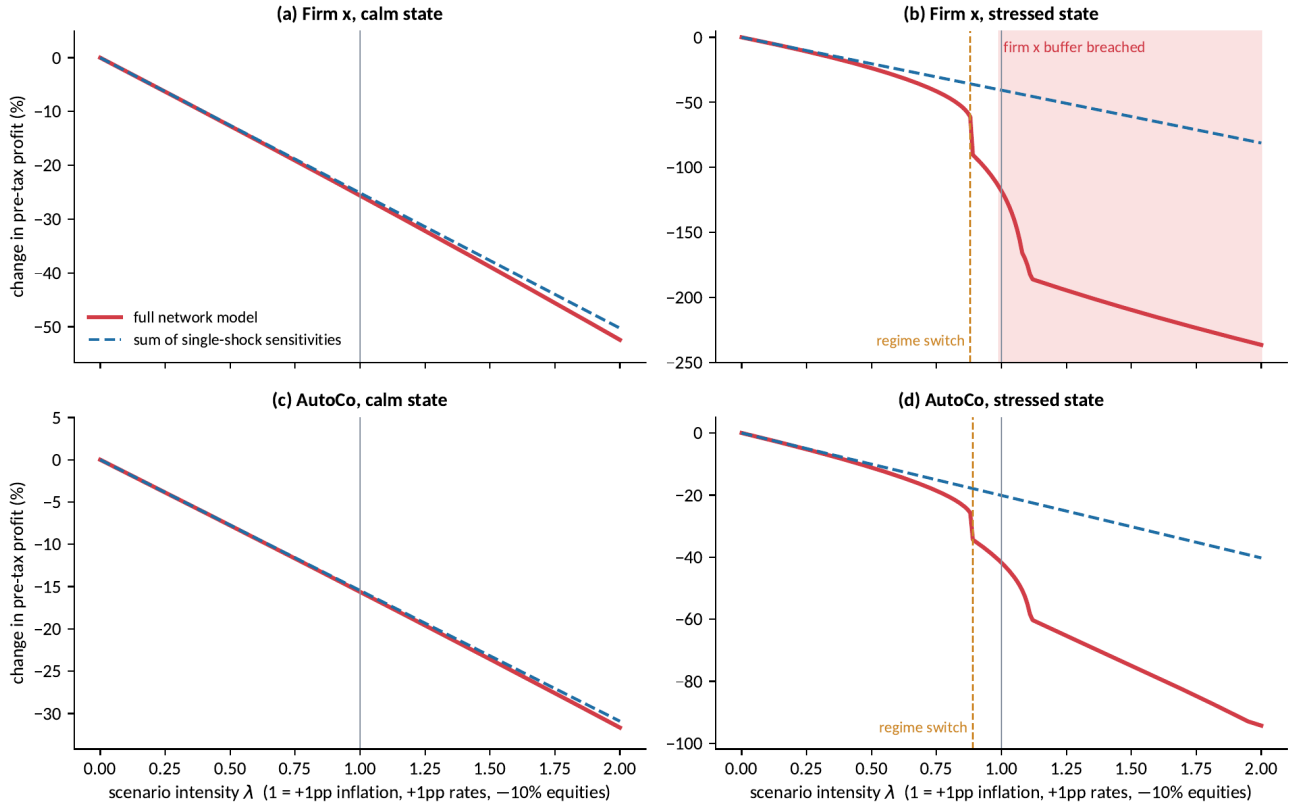


Figure 5. Change in pre-tax profit of firm x and AutoCo as a function of scenario intensity λ, full model against the sum of single-shock sensitivities. Calm state (a, c): the two coincide. Stressed state (b, d): the network switches regime at $\lambda \approx 0.89$, when substitutability collapses; a second step at $\lambda \approx 1.1$ follows from SME failures. Firm x breaches its liquidity buffer at $\lambda \approx 0.98$.

Table 6 and Figures 5b and 5d show the result. The sensitivities measured in the stressed state are larger than in the calm state (by 62% for firm x and 30% for AutoCo), which is a first warning, but they no longer describe the outcome. Up to $\lambda \approx 0.5$ the full model stays close to the expansion; then it bends away, and at $\lambda \approx 0.89$ it jumps. At the full scenario, $\lambda = 1$, firm x loses €4.58 m of pre-tax profit, turning a profit of €3.88 m into a loss of €0.70 m that exceeds its remaining liquidity of €0.5 m. AutoCo loses €4.58 bn, 42% of its pre-tax profit, against 20% from its sensitivities. Its sales volume falls by 11.2% instead of 3.4% in the calm state, the SME periphery loses 16.7% of its volume, and 4.8% of SMEs fail, a third of the vulnerable group. The same four mechanisms of Section 3.2 are at work, each visible in the numbers.

(a) Amplification near criticality. By the spectral decomposition of Eq. (11), applied to G with right eigenvectors u_q, as the largest eigenvalue approaches one a single mode dominates, and every shock is projected onto the same direction u_1. In the stressed state that direction is spread over the whole network (weights 0.13 on AutoCo and ChemCo, 0.15 on LogiCo, 0.26 on the SME periphery and 0.33 on firm x): the loop that runs from AutoCo to its SME suppliers and back dominates. For AutoCo, the share of its own first-round loss that returns to it through the network (the diagonal element of $(I - G)^{-1}$ minus one) rises from 4.8% in the calm state to 30% in the stressed state. The hub is exposed to the echo of its own shock.

(b) State-dependent coupling. As the scenario unfolds, system stress u rises to 10.8%, substitutability falls from 0.50 to its floor of 0.10, and the loop gain of G rises from 0.59 to 0.79. The local loop gain ρ_T, which includes this erosion, climbs from 0.59 to one at $\lambda \approx 0.89$ (Figure 6a): there, the expansion parameter κ of Eq. (7) reaches one, and the solution jumps from the low-stress to the high-stress branch. In the calm state the local loop gain never exceeds 0.45. Because the coupling depends on the shock, shocks interact, and by Eq. (10) interaction terms grow faster than first-order terms as $\rho_0 \to 1$. In the example, a joint inflation and rate shock of 1 pp each costs AutoCo €359 m more than the two shocks separately, against €2 m in the calm state (Figure 6b).

(c) Thresholds: buffers and credit. Thresholds make the response non-differentiable, so no Taylor expansion exists across them. As the SME periphery's margin falls, vulnerable SMEs fail and are not replaced, which cuts supplies to the club; losses then cascade, as in models of default contagion [22,23,24,25]. The credit channel adds an interaction between shocks: inflation lowers EBIT, rates raise interest, and both lower interest cover, so spreads widen by 1.3 pp for AutoCo and by about 6 pp for the SME periphery and firm x.

(d) Dependence between the shocks themselves. As argued in Section 3.2, the joint distribution of inflation, rate and equity shocks depends on the regime. In a stressed state, a large joint move such as the scenario of Table 6 is more likely than a calm-state covariance matrix suggests, and it is exactly the move that the network amplifies most.

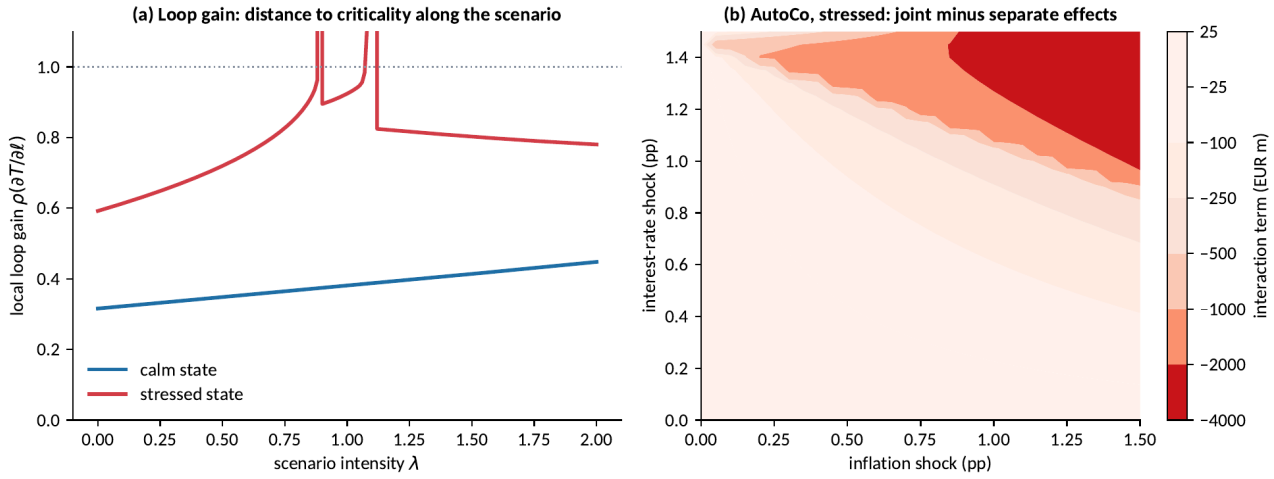


Figure 6. (a) Local loop gain ρ_T, the spectral radius of the Jacobian of the loss map, along the scenario. In the stressed state it reaches one at the two regime switches ($\lambda \approx 0.89$ and $\lambda \approx 1.1$). (b) Interaction between inflation and interest-rate shocks for AutoCo in the stressed state: joint effect minus the two separate effects.

As in Section 3.2, all four properties of Section 3.1 fail. Superposition fails because of interactions, attribution fails because the residual (€2.4 bn for AutoCo) exceeds the attributed part, aggregation fails because a single covariance matrix misses the change in dependence, and hedges sized on Δ_i misfire because Δ_i itself moves. The two firms fail differently. The small firm breaks because of concentration: its fate is set by the hub it supplies and by a buffer it cannot rebuild quickly. The hub breaks because of the network: it is diversified in its sales, but it depends on the periphery collectively for supplies, and when the periphery weakens the stress returns to the hub. Neither could have seen this from its own sensitivities.

E.8 A real example: the 2022 European energy and inflation crisis

The scenario of this paper is not hypothetical. In 2022 inflation rose, interest rates rose and equity markets fell, together, and the shock started in an input with low short-run substitutability: natural gas. Euro area inflation rose from 2.6% in 2021 to 8.4% on average in 2022, with a peak of 10.6% in October [31,32]. The ECB raised its deposit rate from -0.50% in July 2022 to 2.00% in December and 4.00% by September 2023 [33]. The Euro Stoxx 50 fell 23.4% in the first three quarters of 2022 [34]. The Dutch TTF gas price closed at a record of about €339/MWh in late August 2022 and was back near €35/MWh a year later [35]. In the units of Appendix E.3 this was a shock of 5.8 units of inflation (average 2022 against 2021), 2.5 units of interest rates (July to December 2022) and 2.3 units of equities, with a much larger multiple for energy prices (Figure 7a). It was several times larger than the scenario for which our stressed example network switches regime.

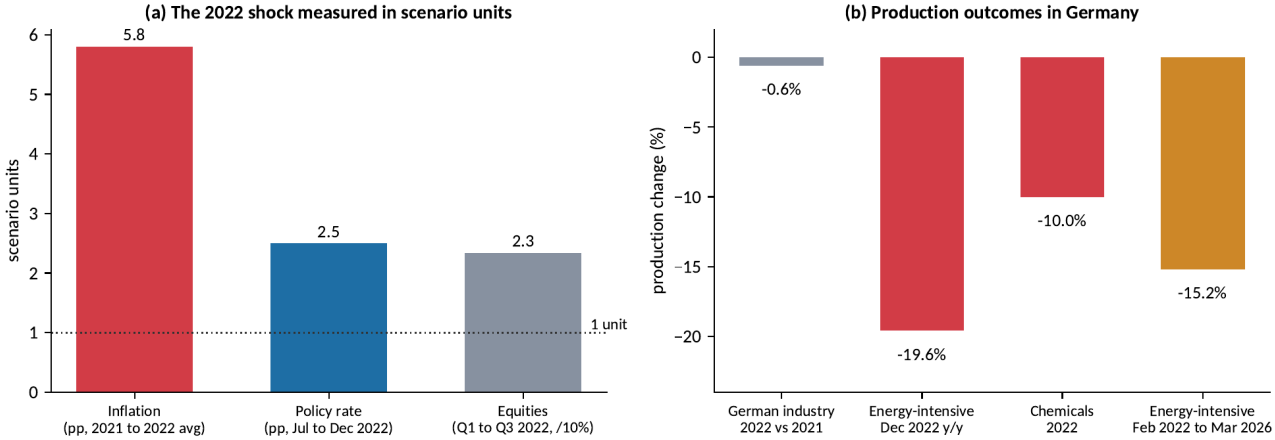


Figure 7. (a) The 2022 shock in scenario units: euro area inflation 2021 to 2022 [31], ECB deposit rate July to December 2022 [33], Euro Stoxx 50 in Q1–Q3 2022 [34]. (b) Production outcomes in Germany: total industry in 2022 [36], energy-intensive industries in December 2022 against a year earlier [36], chemicals in 2022 [37], and energy-intensive industries from February 2022 to March 2026 [38].

What happened. The economy as a whole did not switch regime. German industrial production was only 0.6% lower in 2022 than in 2021 [36], and EU gas consumption fell by 17.7% between August 2022 and March 2023 against its five-year average [39], which shows that substitution and demand reduction worked at the scale of the whole network: in the language of Eq. (12), system-wide substitutability stayed high. The energy-intensive part of the network did switch. Production in German energy-intensive industries was 19.6% lower in December 2022 than a year earlier [36]; chemical production fell 10% in 2022, petrochemicals 15.5%, and one in four chemical companies was loss-making [37]. For those firms, gas was essential (χ close to one) and short-run substitutability was close to zero, margins were thin, and buffers were used up.

The signature of a regime switch. Under perturbation theory the response to a shock is reversible: when the shock goes away, so does the effect. The gas price fell by about 90% from its peak within a year, yet production in energy-intensive industries was

still 15.2% below its February 2022 level in March 2026, against 9.5% for industry as a whole [38]. The subgraph did not return to its former branch when the shock was reversed. This persistence (hysteresis) is what the model of Section 3.2 predicts once a subgraph has crossed its critical point: capacity that was closed, and customers that found other suppliers, do not come back when prices normalise.

Assessment. In the framework of this paper, the 2022 crisis took Europe's energy-intensive subgraph beyond criticality while the wider network remained below it, because substitution at the system level kept the loop gain away from one. Two lessons follow for today. First, the shock size at which a network switches depends on its state, not only on the shock: the same gas shock would have been absorbed by a network with more alternatives and fuller buffers. Second, the energy-intensive subgraph now sits on its high-stress branch with depleted buffers, so its distance to criticality for the next shock is smaller than before 2022.

A numerical simulation: perturbation theory versus active inference. The 2022 episode also gives a clean test of the two approaches of this paper, because the outcome is already known. We compare perturbation theory with active inference by illustration. the energy-intensive subgraph's exposure to the 2022 gas is shown once using perturbation theory in its usual, sensitivity-based form, and once using active inference.

The shock path we fed into both methods (Figure 8a) tracks the shape of the real gas-price move described above: a sharp rise through 2022 and a slower decline back toward a lower, settled level over the following years. Perturbation theory, run on this path, does what the theory is built to do: it reports a response that scales with the shock and relaxes as the shock relaxes, so that once the gas price had fallen by around 90% from its peak, the simulated effect on production had very largely unwound with it (Figure 8b, dashed line). That is a reasonable description of the network as a whole, which did not switch regime, but it is not a reasonable description of the energy-intensive subgraph, which did.

The active-inference simulation was given the same shock path, but allowed the subgraph's simulated state to cross a threshold and change branch, in the manner of Section 3.2, and allowed simulated firms to take actions once that threshold was crossed rather than simply absorbing the shock passively. Two things follow once the threshold is crossed. First, the simulated response is no longer proportional to the shock: it overshoots the shock's own path on the way down, consistent with network effects, not only the direct effect, driving the outcome. Second, and more importantly, the simulated response does not retrace the shock's path back down: it decays far more slowly than the shock itself, leaving a persistent gap even once the shock has largely faded (Figure 8b, solid line). This is the simulated counterpart of the hysteresis described qualitatively above, and it tracks the two real data points plotted alongside it — production in energy-intensive industries around 19.6% lower in December 2022, and still around 15.2% lower in March 2026 — far more closely than the perturbation-theory path does at either date.

The same run also tracks a simulated measure of firm action intensity: the extent to which simulated firms in the subgraph are hoarding inputs, seeking substitutes and tightening terms with counterparties, rather than continuing to behave as they did in the calm state (Figure 8c). In the simulation this action intensity rises sharply as the subgraph approaches and crosses its threshold, stays elevated for an extended period afterward, and decays only slowly — tracking the persistence of the production effect in panel (b) rather than the shock in panel (a). This is the qualitative content of firms "taking actions under criticality": once a subgraph is close to or beyond its critical point, firms are no longer passive recipients of a shock whose effect can be read off a fixed sensitivity; their own responses become part of what is being simulated, and it is those responses that make the effect persist.

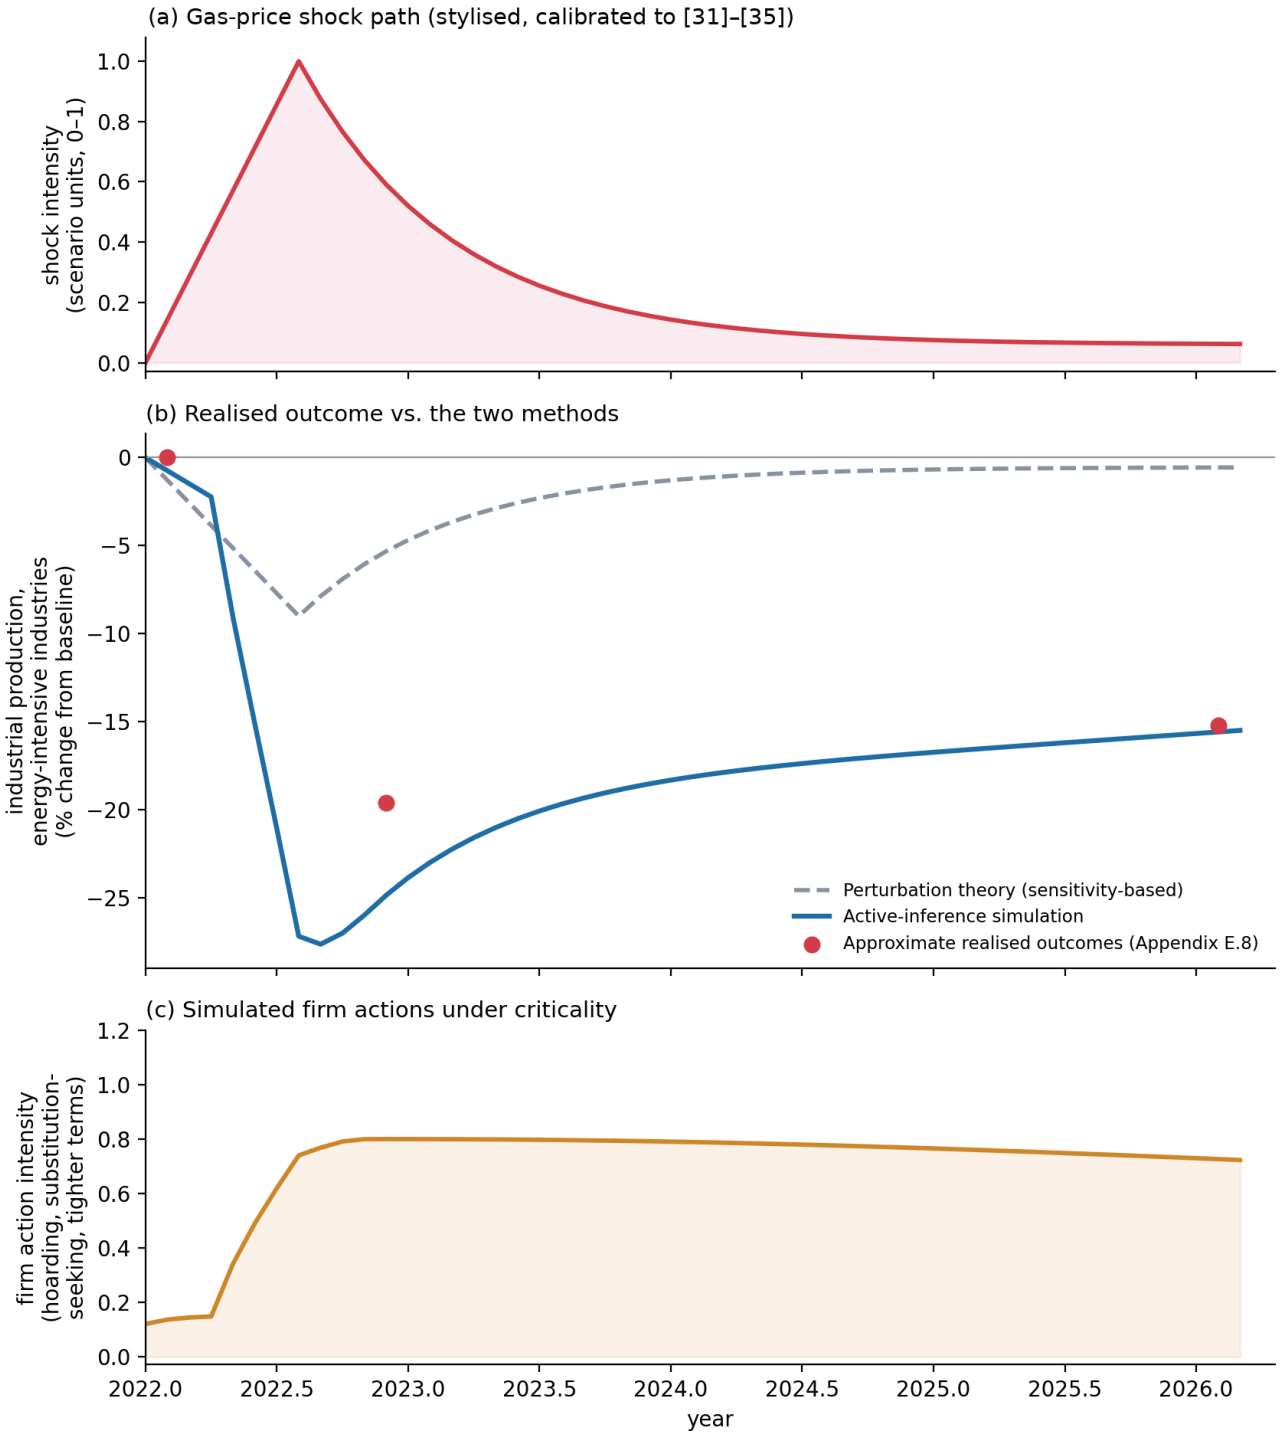


Figure 8. An illustrative numerical simulation of the 2022 episode, built for this paper and calibrated loosely to the data cited above; a toy model, not the paper's underlying framework, and its equations are not reproduced here. (a) Stylised gas-price shock path used as the common input to both methods. (b) The same shock run through perturbation theory (dashed) and through an active-inference-style simulation (solid), against the two real data points discussed above (markers): perturbation theory relaxes with the shock and captures neither the overshoot nor the persistence; the active-inference simulation captures both. (c) Simulated intensity of firm actions (hoarding, substitution-seeking, tighter terms) in the active-inference run, which rises as the subgraph crosses its threshold and decays only slowly thereafter, in step with the persistent effect in (b) rather than with the shock in (a).

This does not show that any particular active-inference model will match a given crisis; it illustrates, on a case where the outcome is already known, the qualitative gap between the two approaches that the rest of this paper argues for analytically. A sensitivity-based method is built to be reversible and so cannot represent a regime switch by construction, however it is calibrated. A method that lets the network's state, and firms' own actions, feed back into the simulation can represent regime switch include the hysteresis effect when shocks reduce.

References

1. De Jonckheere, T., Putseys, T. (2026). European Resilience: Why European resilience requires autonomy outside and interdependence inside. hedgeo white paper, September 2026.

2. Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica* 80(5), 1977–2016. https://doi.org/10.3982/ECTA9623
3. Carvalho, V.M., Tahbaz-Salehi, A. (2019). Production networks: A primer. *Annual Review of Economics* 11, 635–663.
4. Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica* 79(3), 733–772. https://doi.org/10.3982/ECTA8769
5. Bacilieri, A., Borsos, A., Astudillo-Estévez, P., Hoefer, M., Lafond, F. (2026). Firm-level production networks: what do we (really) know? *Journal of Economic Dynamics and Control* 187, forthcoming. Working paper version: INET Oxford Working Paper No. 2025-14. https://ideas.repec.org/p/amz/wpaper/2023-08.html
6. The Register (2025). Jaguar Land Rover gets £1.5bn government jump-start after cyber breakdown. 29 September 2025. https://www.theregister.com/2025/09/29/jlr_government_loan/
7. CBT News (2025). Global auto production at risk as Nexperia chip halt sparks new supply chain crisis. 29 October 2025. https://www.cbtnews.com/global-auto-production-at-risk-as-nexperia-chip-halt-sparks-new-supply-chain-crisis/
8. Carvalho, V.M., Nirei, M., Saito, Y.U., Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics* 136(2), 1255–1321. https://doi.org/10.1093/qje/qjaa044
9. Barrot, J.-N., Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics* 131(3), 1543–1592. https://doi.org/10.1093/qje/qjw018
10. Atalay, E., Hortacsu, A., Roberts, J., Syverson, C. (2011). Network structure of production. *Proceedings of the National Academy of Sciences* 108(13), 5199–5202. https://doi.org/10.1073/pnas.1015564108
11. Bernard, A.B., Moxnes, A., Saito, Y.U. (2019). Production networks, geography, and firm performance. *Journal of Political Economy* 127(2), 639–688. https://doi.org/10.1086/700764
12. Barabasi, A.-L., Albert, R. (1999). Emergence of scaling in random networks. *Science* 286(5439), 509–512. https://doi.org/10.1126/science.286.5439.509
13. Zhou, S., Mondragon, R.J. (2004). The rich-club phenomenon in the Internet topology. *IEEE Communications Letters* 8(3), 180–182.
14. Colizza, V., Flammini, A., Serrano, M.A., Vespignani, A. (2006). Detecting rich-club ordering in complex networks. *Nature Physics* 2, 110–115. https://doi.org/10.1038/nphys209
15. Diem, C., Borsos, A., Reisch, T., Kertész, J., Thurner, S. (2022). Quantifying firm-level economic systemic risk from nation-wide supply networks. *Scientific Reports* 12, 7719. https://doi.org/10.1038/s41598-022-11522-z
16. Baqaee, D.R., Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica* 87(4), 1155–1203. https://doi.org/10.3982/ECTA15202
17. Elliott, M., Golub, B., Leduc, M.V. (2022). Supply network formation and fragility. *American Economic Review* 112(8), 2701–2747. https://doi.org/10.1257/aer.20210220
18. Tabachová, Z., Diem, C., Borsos, A., Burger, C., Thurner, S. (2024). Estimating the impact of supply chain network contagion on financial stability. *Journal of Financial Stability* 75, 101336. https://doi.org/10.1016/j.jfs.2024.101336
19. Diem, C., Borsos, A., Reisch, T., Kertész, J., Thurner, S. (2024). Estimating the loss of economic predictability from aggregating firm-level production networks. *PNAS Nexus* 3(2), pgae064. https://doi.org/10.1093/pnasnexus/pgae064
20. Leontief, W.W. (1936). Quantitative input and output relations in the economic systems of the United States. *Review of Economics and Statistics* 18(3), 105–125.
21. Horn, R.A., Johnson, C.R. (2013). *Matrix Analysis*, 2nd edn. Cambridge University Press.
22. Eisenberg, L., Noe, T.H. (2001). Systemic risk in financial systems. *Management Science* 47(2), 236–249. https://doi.org/10.1287/mnsc.47.2.236.9835
23. Battiston, S., Puliga, M., Kaushik, R., Tasca, P., Caldarelli, G. (2012). DebtRank: Too central to fail? Financial networks, the FED and systemic risk. *Scientific Reports* 2, 541. https://doi.org/10.1038/srep00541
24. Haldane, A.G., May, R.M. (2011). Systemic risk in banking ecosystems. *Nature* 469, 351–355. https://doi.org/10.1038/nature09659
25. Buldyrev, S.V., Parshani, R., Paul, G., Stanley, H.E., Havlin, S. (2010). Catastrophic cascade of failures in interdependent networks. *Nature* 464, 1025–1028. https://doi.org/10.1038/nature08932
26. Campbell, J.Y., Pflueger, C., Viceira, L.M. (2020). Macroeconomic drivers of bond and equity risks. *Journal of Political Economy* 128(8), 3148–3185. https://doi.org/10.1086/707766

27. Longin, F., Solnik, B. (2001). Extreme correlation of international equity markets. *Journal of Finance* 56(2), 649–676. https://doi.org/10.1111/0022-1082.00340

28. Ang, A., Chen, J. (2002). Asymmetric correlations of equity portfolios.
Journal of Financial Economics

63(3), 443–494.

29. Hamilton, J.D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57(2), 357–384. https://doi.org/10.2307/1912559

30. Scheffer, M., Bascompte, J., Brock, W.A., Brovkin, V., Carpenter, S.R., Dakos, V., Held, H., van Nes, E.H., Rietkerk, M., Sugihara, G. (2009). Early-warning signals for critical transitions. *Nature* 461, 53–59. https://doi.org/10.1038/nature08227 %BIBEND

31. Eurostat. Inflation in the euro area. Statistics Explained. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Inflation_in_the_euro_area

32. Eurostat (2022). Annual inflation up to 10.6% in the euro area. Euro indicators, 17 November 2022. https://ec.europa.eu/eurostat/documents/2995521/15265521/2-17112022-AP-EN.pdf

33. European Central Bank (2023). Annual Report 2022. https://www.ecb.europa.eu/pub/pdf/annrep/ecb.ar2022~8ae51d163b.en.pdf

34. Stock Market Walk (2022). Euro Stoxx 50 performance: what happened in 2022? 3 October 2022. https://stockmarketwalk.com/stock-market/indices/euro-stoxx-50-performance/6622/

35. Euronews (2023). The big turnaround: how Europe's gas prices fell from €300 to €35 MWh in the span of a year. 28 August 2023. https://www.euronews.com/my-europe/2023/08/28/the-big-turnaround-how-europes-gas-prices-fell-from-300-to-35-mwh-in-the-span-of-a-year

36. Destatis (2023). Production in December 2022: -3.1 on the previous month. Press release, February 2023. https://www.destatis.de/EN/Press/2023/02/PE23_049_421.html

37. ICIS (2022). German chems, pharma output drops 6% in 2022, prospects bleak – VCI. 15 December 2022. https://www.icis.com/explore/resources/news/2022/12/15/10835857/german-chems-pharma-output-drops-6-in-2022-prospects-bleak-vci/

38. Clean Energy Wire (2026). Energy-intensive industries in Germany see 15-percent production fall since start of Ukraine war (Destatis data). 15 May 2026. https://www.cleanenergywire.org/news/energy-intensive-industries-germany-see-15-percent-production-fall-start-ukraine-war

39. Eurostat (2023). EU gas consumption, August 2022 to March 2023, against the 2017–2022 average. News item, 19 April 2023. https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20230419-1