


BEYOND PROCUREMENT'S PROBLEM

Supply Chain Fragility: The Programme Risk Nobody Modelled

A Practitioner's Perspective on Supply Chain as a First-Class Programme
Dependency

Giovanni Leonardi

April 2020



Contents

1. Executive Summary	3
2. How a Whole-Programme Risk Became Nobody's Job	3
3. What the Evidence Now Shows	4
4. Why Programmes Are Especially Exposed	5
5. A Model for Treating Supply as a First-Class Programme Risk	6
5.1 Map the dependency, not just the supplier	6
5.2 Stress the chain, do not merely list it	6
5.3 Buffer deliberately, and price it honestly	6
5.4 Govern it at the programme table	7
6. Recommendation	7

Supply Chain Fragility: The Programme Risk Nobody Modelled

Programme · No. 188

Originally written April 2020 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

"A dependency that can halt every workstream, yet appears on no programme register, is not a low risk — it is an unmeasured one."

1. Executive Summary

Programme management has spent two decades refining the risks it takes seriously. Cost, schedule, scope, quality, benefits realisation, stakeholder alignment, technology delivery — each now has its own vocabulary, its own governance forum, its own line on the register. One dependency, however, has remained stubbornly invisible: the physical supply chain on which a programme's outcomes ultimately rest. It was treated as procurement's concern, or operations', or somebody else's — but rarely the programme's own. The first months of this year have exposed the price of that blind spot.

The argument of this paper is straightforward. Supply chain fragility is not an operational nuisance to be absorbed downstream; it is a first-class programme risk that belongs on the register from the outset, modelled with the same rigour we apply to critical path and cost contingency. When a single supplier, a single port, or a single geography can arrest every workstream in a programme simultaneously, the exposure is not marginal. It is systemic. And systemic exposure that appears nowhere in the risk model is not a low risk — it is an unmeasured one.

What follows sets out why the omission happened, what the evidence of recent weeks reveals about its consequences, and a defensible model for putting supply dependency where it belongs: at the centre of programme governance rather than the edge of it.

2. How a Whole-Programme Risk Became Nobody's Job

The roots of the blind spot are structural, not careless. Three forces combined over the last twenty years to push supply chain risk off the programme radar.

The first was the triumph of lean thinking. Just-in-time inventory, pioneered in automotive manufacturing and then generalised across nearly every sector, taught a generation of managers that inventory is waste, that buffers are inefficiency made visible, and that the tightest supply chain is the best one. This was, on its own terms, correct: working capital tied up in warehouses earns nothing, and slack invites complacency. But

lean optimised relentlessly for cost and flow under normal conditions, and in doing so it quietly traded away resilience. The buffer that looked like waste on a spreadsheet was, in a different light, the programme's insurance policy.

The second force was globalisation of sourcing. Over two decades the supplier base for almost everything — components, chemicals, packaging, active pharmaceutical ingredients, finished goods — consolidated into a smaller number of lower-cost geographies, often a single one. The efficiency case was overwhelming and the risk case was abstract, so the risk case lost. Concentration that would have alarmed a risk committee if it appeared in a credit portfolio was treated as ordinary good practice when it appeared in a bill of materials.

The third force was organisational. Programmes are structured around delivery: the workstreams, the milestones, the benefits case. Procurement and supply sit in a different reporting line, speak a different language, and are engaged transactionally — issue the purchase order, manage the contract, escalate if something slips. The result is that the people who understand the supply chain are rarely in the room where programme risk is discussed, and the people in that room treat supply as a solved problem sitting one layer down.

The buffer that lean thinking taught us to see as waste was, in a different light, the programme's insurance policy. We did not remove the risk when we removed the inventory. We simply stopped paying the premium and hoped the claim would never come.

Each of these forces was individually rational. Together they produced an outcome no one intended: a category of risk capable of stopping an entire programme, owned by no one who sat at the programme table.

3. What the Evidence Now Shows

The events of recent weeks have provided a natural experiment that no risk workshop could have designed. Several patterns are already unmistakable.

- Single-geography concentration converted a regional event into a global stoppage. When manufacturing in one region halted early in the year, the effect did not stay regional. It propagated along dependency chains that most organisations had never fully mapped, surfacing in finished-goods shortages weeks later and thousands of miles away. Programmes discovered dependencies they did not know they held.
- The bullwhip effect returned at full force. Small shifts in end demand — panic purchasing in some categories, collapse in others — amplified violently upstream. Ordering signals that swung wildly left suppliers unable to plan, and the resulting whiplash disrupted supply even where underlying capacity was intact.
- Logistics, not just production, proved to be the fragile link. The collapse of passenger aviation removed a vast share of air-freight belly capacity almost overnight, stranding shipments that were manufactured and ready but could not move. Programmes that had modelled supplier risk had almost universally failed to model

the freight path.

- Tier-two and tier-three suppliers were the real exposure. Organisations had visibility of their direct suppliers and almost none beyond them. The failures that hurt most often originated two or three tiers down, at a sub-component maker nobody in the programme had ever heard of.

None of these patterns is genuinely new. The Tohoku earthquake and the Thai floods nearly a decade ago taught the same lessons about concentration and hidden tiers, and the literature on the bullwhip effect is older still. What is new is the scale and simultaneity. Previous shocks were geographically bounded; an organisation could source around them. This one arrived everywhere at once, and the workarounds that resilience planning had quietly assumed — shift to an alternate region, expedite by air, draw down a buffer — were unavailable together, at the same moment, to everyone.

“Resilience plans failed not because the individual mitigations were wrong, but because every mitigation assumed the rest of the world was still functioning normally.”

4. Why Programmes Are Especially Exposed

It is worth being precise about why this is a programme problem and not merely a supply-chain problem, because the distinction determines who must act.

An operational supply chain that falters degrades service: shelves are thinner, lead times stretch, some customers are disappointed. Painful, but usually recoverable and rarely existential. A programme is different in three respects that sharpen the exposure considerably.

1. Programmes are time-boxed against a benefits case. A delay that operations can absorb as a bad quarter can, for a programme, breach the very window in which the benefits were to be realised. A six-week supply interruption is an inconvenience to a running operation and potentially fatal to a milestone-driven programme with a hard external deadline.
2. Programmes concentrate dependency at integration points. Workstreams are designed to converge — a physical build, a site rollout, a hardware deployment. A supply failure at a convergence point does not delay one workstream; it stalls the integration, and everything downstream of it, at once.
3. Programmes carry sunk commitment that removes flexibility. By the time a supply risk materialises, the programme has typically committed to a design, a vendor, a specification. The graceful degradation available to an operation — substitute a product, drop a line — is often contractually or technically closed to a programme mid-flight.

This is why a risk that operations can treat as a manageable cost of doing business must, for a programme, be treated as a potential single point of failure. The same disruption lands with entirely different force depending on where it lands.

5. A Model for Treating Supply as a First-Class Programme Risk

Diagnosis without remedy is commentary. The remainder of this paper sets out a practical model for bringing supply dependency onto the programme register and governing it there. It has four components: map, stress, buffer, and govern.

5.1 Map the dependency, not just the supplier

The first discipline is to extend visibility beyond the first tier. For each critical programme deliverable, trace the supply dependency to the point of genuine single-sourcing — the tier at which no realistic alternative exists. This is laborious and it is resisted, because direct suppliers are often reluctant to expose their own sources. But a dependency map that stops at tier one is a map of the part of the risk you can already see. The part that stops programmes lives further down.

The practical output is a critical-path dependency register that names, for each essential input, the geography of true origin, the number of qualified sources, and the switching lead time to an alternative. Where that register shows a single geography, a single source, and a switching time longer than the programme's tolerance, the programme has found a systemic risk — whatever the risk workshop concluded.

5.2 Stress the chain, do not merely list it

A register that records dependencies without testing them offers false comfort. The second discipline is to subject the critical dependencies to deliberate stress scenarios: what happens to the milestone if this geography closes for eight weeks; if freight capacity on this lane halves; if this sole supplier fails entirely. The purpose is not precise prediction — that is unavailable — but to convert an abstract dependency into a concrete, quantified impact on the plan that governance can actually weigh.

Scenario	What it tests	Typical blind spot exposed
Single-geography closure	Concentration risk	Hidden lower-tier co-location
Sole-supplier failure	Substitution readiness	Absence of qualified alternates
Freight-lane loss	Logistics path	Supplier modelled, route ignored
Demand whiplash	Ordering discipline	Amplification up the chain

5.3 Buffer deliberately, and price it honestly

The third discipline is to restore buffering as a conscious choice rather than an embarrassment. Lean was right that undisciplined inventory is waste; it was wrong to conclude that all buffering is. For a programme, a strategic buffer — of critical

components, of qualified alternate suppliers pre-contracted, of scheduled float at integration points — is not inefficiency. It is contingency in physical form, and it should be costed and approved exactly as financial contingency is. The governing question is not how do we eliminate this buffer but what shock is this buffer insuring against, and is the premium proportionate to the exposure.

5.4 Govern it at the programme table

The final discipline is ownership. Supply dependency must have a named owner who sits at the programme risk forum, not one layer below it, and the critical-path dependency register must be a standing item alongside cost and schedule. The test of whether an organisation has absorbed the lesson of these weeks is simple: at the next programme board, is supply concentration discussed with the same seriousness as budget contingency? If it is raised only when something has already failed, nothing structural has changed.

6. Recommendation

The recommendation of this paper is specific and deliberately narrow, because a narrow change that is actually adopted beats a broad one that is admired and shelved.

Every programme should carry a critical-path supply dependency as a named, owned, and stress-tested entry on its risk register, reviewed at the programme board with the same standing as cost and schedule contingency. Concretely, that means three things: trace each essential input to its true single-source tier and geography; stress the resulting dependencies against a small set of severe-but-plausible scenarios; and hold a strategic buffer wherever the stress test shows a switching time longer than the programme can tolerate, priced and approved as contingency.

None of this requires new methodology. It requires only that we stop treating the physical supply chain as somebody else's operational detail and start treating it as what the evidence now plainly shows it to be — a dependency capable of halting the whole endeavour, and therefore a programme risk of the first order. The organisations that will emerge from this period strongest will not be those that were luckiest with their suppliers. They will be those that decide, having seen the pattern once, never to leave it unmodelled again.