

BEYOND THE BUFFER

Resilience Versus Efficiency Is the Wrong Argument

What eighteen months of broken supply chains revealed about how manufacturers actually value resilience — and why the current enthusiasm for it will not last.

Giovanni Leonardi

July 2021



Resilience Versus Efficiency Is the Wrong Argument

Transformation · No. 265

Originally written July 2021 | Re-edited and re-rendered for the WhereBeyond Edition, 2025

“The parts that can stop you are rarely the parts that cost you, and our reporting is built entirely around the second.”

1. The Part That Cost Pennies

One story stood in for the whole year. A production line worth several million pounds a week, brought to a halt — not by a semiconductor, not by a rare-earth metal, but by a moulded rubber seal costing a few pence. One qualified supplier. One factory. One shipping lane. When the factory closed for a fortnight, the line stood idle, and the finished product — worth thousands of times the value of the missing part — could not be built.

Everyone in operations has now collected a version of this story. The specifics change — a wiring harness, a specialty resin, a particular grade of foam, a controller chip — but the shape is identical. A tiny, cheap, invisible component, sourced from a single point, sitting on the critical path of something enormous. For a decade we optimised these components out of sight. This year they optimised us.

Out of that experience a slogan has hardened: resilience versus efficiency. Boards now ask for resilience the way they once asked for cost-out. The framing has settled into a comfortable trade-off — a dial to be turned, a little away from efficiency, a little towards resilience, and the only debate is how far. I want to argue that this framing is wrong, and quietly dangerous.

2. The Trade-Off That Was Never Priced

The trouble with resilience versus efficiency is that it dignifies as a choice something most organisations never actually chose. A trade-off implies two quantities on the same scale, weighed against each other. But we never had two quantities. We had one number we measured obsessively — landed cost, working capital, inventory turns — and another we did not measure at all.

Consider what actually happened over the last decade of sourcing decisions. Consolidating two suppliers into one saved, say, three per cent on unit cost. That three per cent was real, visible, bankable, and it landed in someone's objectives that year. The risk it created — a single point of failure on the critical path — was invisible. It carried no line on any report. It cost nothing until the year it cost everything.

We did not weigh resilience against efficiency and choose efficiency. We measured efficiency, left resilience unmeasured, and mistook the silence for a zero.

This is not a trade-off. It is a measurement asymmetry. And it matters enormously which of the two we are dealing with, because the remedies are opposite. If the problem is a genuine trade-off, the answer is to re-balance — hold more stock, dual-source more parts, accept a permanently higher cost base. If the problem is a measurement failure, adding cost everywhere is exactly the wrong move: it is paying blindly to fix something we still cannot see.

3. Why "More Buffer" Is the Wrong Lesson

The instinct of the past few months has been to buffer everything. Inventory targets up across the board. A second supplier for everything. Reshoring studies commissioned by the dozen. Safety stock rebuilt to levels not seen since before the lean decades. Some of this is sensible. Most of it is the last war being fought with a blunt instrument.

Here is the steelman for the efficiency position, and it deserves to be taken seriously rather than dismissed as recklessness. Lean did not happen by accident or by stupidity. Carrying inventory ties up cash, hides quality problems, ages into obsolescence, and fills warehouses with product that customers may never want. Dual-sourcing splits volume, weakens negotiating leverage, doubles qualification and audit cost, and multiplies the number of relationships that can go wrong. For most parts, in most years, the resilience premium is pure waste — insurance against an event that does not come. The people who built just-in-time were not fools; they were responding, correctly, to decades in which the buffers really were fat and the disruptions really were rare. We were warned otherwise — the earthquake and the floods of 2011 stopped lines the same way this year has — and we chose, reasonably enough at the time, to treat those as freak events rather than as the standing condition. A generation of genuine value came from taking that slack out.

So the zealot's answer — buffer everything, reshore everything, never single-source again — is not the wisdom of this year. It is the same undisciplined thinking that got us here, run in reverse. Optimising blindly for cost gave us brittle chains. Optimising blindly for resilience will give us bloated ones, and the moment margins tighten — and they always tighten — the blunt buffers will be the first thing cut, by the very logic that stripped them out before. We will have paid for resilience and kept none of it.

4. Resilience as a Design Property, Not a Quantity of Stock

The way out is to stop treating resilience as an amount and start treating it as a design property — something engineered into specific places for specific reasons, and therefore something you can defend when the cost pressure returns. That means doing the unglamorous work most organisations skipped.

- Map the critical path, not the spend. Procurement knows exactly where the money is. Almost nobody knows where the fragility is. The rubber seal does not appear on a spend-ranked supplier list; it is a rounding error. The parts that can stop you are rarely the parts that cost you, and our reporting is built entirely around the second.
- Know the second and third tiers. The disruption this year did not mostly happen at Tier 1 — it happened two and three levels down, at sub-suppliers our direct suppliers themselves could not always name. Many organisations discovered, in the middle of a shortage, that they had no idea who actually made the thing inside the thing they bought.
- Price the failure, not just the part. A single-source decision that saves three per cent looks brilliant until you put a number on the fortnight of lost production it risks. When the line-down cost is eight weeks of output and the saving was three per cent of a cheap component, the efficient choice was never efficient. It was unpriced.
- Place resilience selectively. Not every part deserves a second source; most do not. The discipline is to spend the resilience budget where a failure is both plausible and catastrophic, and to leave the rest genuinely lean — so that when the cost review comes, you can say precisely why each buffer exists and what it protects.

The difference between this and "hold more stock" is the difference between a decision you can defend and one you cannot. Blanket buffers have no owner and no rationale; they are cut the moment someone needs a number to hit. A specific buffer, protecting a specific catastrophic failure point, with a specific price on the failure it prevents, survives the next efficiency drive — because it can explain itself.

5. The Number That Should Change

If only one thing changes as a result of this year, it should be this: the single-source, single-site, single-lane decision should never again be allowed to book its saving without booking its risk in the same breath.

Today a sourcing manager who consolidates to one supplier is credited with the saving immediately and charged for the failure never — or charged only years later, when a different person is in the seat and the disruption is treated as an act of God. That asymmetry is the whole problem in miniature. It is why the chains got brittle, and it is why they will get brittle again once this year's fear fades, unless the accounting changes.

The mechanism need not be elaborate. A single-source or single-site decision above a materiality threshold carries an estimated exposure — days of production at risk, multiplied by the daily contribution margin of what the line produces — recorded next to the saving and reviewed by the same governance that approved it. It will be a rough number, and it does not need to be precise to be transformative. Its whole purpose is to end the era in which the saving had a figure and the risk had none, so that a three per cent cost reduction that puts eight weeks of margin on a single supplier's roof can no longer be presented as an unqualified win. Make the exposure visible and priced, and most of the reckless single-sourcing simply stops being approved — not because anyone forbids it, but because it no longer looks free.

“Resilience will not survive as a slogan. It survives only if the risk shows up on the same page as the saving, in the same year, against the same person's name.”

The organisations that come out of this stronger will not be the ones that bought the most inventory. They will be the ones that changed what they measure — that made the invisible risk visible, put a price on it, and pushed that price back to the moment and the person making the decision. Everyone else is buying buffers they will not keep, congratulating themselves on a resilience they have not designed, and waiting — though they do not yet know it — to relearn the same lesson the next time a part that costs pennies stops a line that earns millions.

That is the real choice this year offered. Not resilience or efficiency. A supply chain that knows where it is fragile, or one that finds out the hard way, again.