

BEYOND THE EFFICIENCY DIAL

Resilience Was Never the Enemy of Efficiency

Why the past eighteen months exposed a measurement failure, not a trade-off

Giovanni Leonardi

July 2021



Resilience Was Never the Enemy of Efficiency

Transformation · No. 264

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

"The saving was a rounding error dressed as a triumph."

1. The connector that stopped the line

A line stops. Not for want of a machine, or a skilled operator, or a customer order — the orders are stacked higher than the plant has seen in two years. It stops for want of a moulded connector that costs less than a cup of coffee, made in one building, in one town, that has gone quiet. The plant manager can tell you to four decimal places what that connector costs. What he cannot tell you — what nobody in the building can tell you — is who else makes it, how many days of stock remain, or how long it would take to qualify an alternative. The price was known to the fraction of a penny. The exposure was not known at all.

That scene, in one variation or another, has played out across thousands of organisations these past eighteen months. A pandemic emptied shelves and then, more strangely, refilled order books faster than factories could answer. A single grounded vessel closed a canal for the better part of a week and rippled through schedules for months. A shortage of semiconductors idled vehicle assembly lines that had nothing wrong with them but the absence of a chip smaller than a fingernail. Boards that had never in their lives asked a question about a tier-two supplier now asked about little else.

And out of all this has come a debate that is now conducted in almost every operations review I sit in: resilience or efficiency? As though the last year and a half had handed us a dial, and our task were simply to turn it — away from the lean, low-cost, tightly-optimised machine we spent two decades building, and towards something safer, slacker, and considerably more expensive. I want to argue that the dial does not exist, that the question is the wrong one, and that treating it as a trade-off is precisely the habit of mind that made us fragile in the first place.

2. The dial that was never real

The trade-off framing is seductive because it flatters everyone. It tells the efficiency camp that they were right all along and merely unlucky — that resilience is a cost you bolt on when the world turns hostile, an insurance premium against a storm nobody could have forecast. And it tells the resilience camp that the solution is simply more: more stock, more suppliers, more slack, a general loosening of the belt that the last decade tightened.

Both camps agree on the shape of the world. They disagree only about which way to turn the dial.

I think both are wrong, and wrong in the same way. Resilience and efficiency are not two ends of a single axis. A well-designed supply chain is not one that has sacrificed some efficiency to buy some resilience; it is one that was never as efficient as its scorecard claimed, because the scorecard was measuring the wrong thing.

The organisations that came through the shocks best were not the ones that had chosen resilience over efficiency. They were the ones that had never mistaken cheap for efficient.

That distinction sounds like wordplay until you look at how the decisions were actually made.

3. What we were actually measuring

For twenty years, the discipline rewarded a particular and very measurable kind of virtue. Unit cost down. Inventory turns up. Days of stock down. Supplier count consolidated. Each of these is a real number, each sits cleanly on a procurement scorecard, and each moves in a direction everyone can celebrate at the quarterly review. I have watched a category manager collect a bonus for a sourcing decision that, three years later, closed a plant for the better part of a season. Both things were true. The saving was real. The exposure was real. Only one of them was ever written down.

Consider the arithmetic that governed a decision I have seen made, in substance, dozens of times. Consolidating a component to a single low-cost source saves nine per cent on that component's unit price. The component is four per cent of the product's bill of materials. So the whole product costs about a third of a per cent less to build — 0.36 per cent, booked in full, banked, and applauded. Then the source goes dark. The line is down for eleven weeks. At that plant a day of lost output runs to something like £120,000 in forgone contribution, and two long-standing customers quietly move volume to a competitor and never fully bring it back. The saving was a rounding error dressed as a triumph. The failure was a catastrophe nobody had costed, because there was no box on the form in which to cost it.

This is the heart of it. What we called efficiency was, very often, local cost minimisation measured on a spreadsheet that had no cell for the cost of failure. The efficiency was real as far as the spreadsheet could see, and the spreadsheet could not see the thing that mattered most.

What the scorecard measured	What it could not see
Unit price of the component	The cost of the line stopping when the component cannot be had
Inventory turns and days of stock	The time it would take to recover once stock ran out

What the scorecard measured	What it could not see
Number of suppliers, consolidated	Whether those suppliers all depended on the same source beneath them
Landed cost per unit, this quarter	The revenue that leaves and does not return when you cannot supply

The single most revealing exercise I have run in the past year takes an afternoon and costs nothing. Ask an organisation to name, for its most critical component, every tier-two supplier sitting beneath its tier-one suppliers. In case after case the visibility stops dead at the first tier. One manufacturer proud of having "diversified" a critical part across four suppliers discovered, when it finally mapped the chain, that all four bought from the same plant two tiers down. The diversification was cosmetic. Four contracts, one point of failure, and a resilience story that fell apart the moment anyone actually looked.

4. The strongest case for lean — and why it does not hold

It would be easy, and lazy, to tell this story as though lean itself were the villain. It is not, and the people who built these systems were not fools. So let me put their case as strongly as I can, because it is a serious one.

For two decades, lean and just-in-time delivered enormous and entirely real value. Capital that would have sat dead in warehouses went to work in the business. Quality improved because defects surfaced immediately instead of hiding in months of buffer stock. Discipline improved because slack conceals problems and lean removes the slack. The shocks of the past year, the argument runs, were once-in-a-century events — and running a business permanently provisioned for the hundred-year storm destroys the very returns that keep it alive. The firms now panic-buying inventory will be sitting on write-downs the moment demand normalises. Most years, the lean operator wins, and wins handsomely. You do not rebuild your entire operating model around the worst eighteen months in living memory.

Every part of that is true, and I would not want to be read as arguing against it. But it contains a quiet substitution that gives the whole game away. The choice was never lean versus bloated. It was never hold no stock versus hold stock everywhere. That was a caricature of lean, and it was the caricature — not the discipline — that stripped the buffers that mattered.

Here is the fact that ought to end the debate. The very manufacturer that gave the world just-in-time — the originator of the whole lean philosophy — was, through the early part of this year, conspicuously less disrupted by the semiconductor shortage than its rivals. Why? Because after an earthquake devastated its home supply base a decade ago, it had done something that looks, superficially, like a betrayal of its own religion: it deliberately held weeks of buffer stock of specific, hard-to-source components. Not stock everywhere. Stock precisely where recovery would be slow and the part irreplaceable. The lesson the rest of us took from lean was "hold no inventory." The lesson its own inventor took from its own history was "hold inventory exactly where it cannot be

replaced quickly." Those are opposite lessons, drawn from the same teacher. We learned the slogan. They understood the design.

5. Resilience as a property of the design

So resilience is not slack added back. It is not the dial turned down from efficiency. It is what you get when you price failure correctly and then let that price shape the design — and, done properly, it removes far more waste than it adds, because it stops you spending precious buffer on the things that never needed protecting.

The organisations doing this well are not holding more of everything. They are holding almost nothing of most things and a great deal of a few things, and they can tell you exactly why. The discipline is to place redundancy where three conditions meet:

- Criticality — the component or supplier whose absence stops the line or the sale, not the one that is merely expensive.
- Fragility of source — the part that comes from a single plant, a single region, or a chain you cannot see the bottom of.
- Slowness of recovery — the part that takes months to re-qualify, re-tool, or re-certify, where money cannot buy back the lost time.

Where all three coincide, buffer stock and second sources are not waste; they are the cheapest insurance you will ever buy. Where none of them holds, buffer stock is exactly the dead capital the lean camp rightly despises, and piling it up in a panic is not resilience — it is anxiety with a warehouse. The skill is not more or less. It is which.

And underneath all of it sits the unglamorous discipline that makes the rest possible: visibility. You cannot design redundancy for a chain you cannot see. Mapping past the first tier is tedious, it is resisted by suppliers who regard their own sources as commercial secrets, and it produces no saving that shows up this quarter. It is also the single highest-return piece of work most organisations could do right now, and almost none of them had done it before they were forced to.

"We spent twenty years optimising the number we could see, and were undone by the number we had chosen not to look at."

6. The equal and opposite mistake

If the framing of this piece has a warning inside it, this is the moment to make it plain — because the trade-off mindset is dangerous in both directions, and the over-correction is already visible.

Having been humbled by fragility, a good many organisations are now doing the mirror image of what got them here: turning the imaginary dial hard the other way. Doubling stock across the board. Adding second and third suppliers indiscriminately. Reshoring on instinct rather than analysis. Writing "resilience" into every objective as though the word itself were a strategy. This feels like learning. It is not. It is the same error inverted — treating resilience as a quantity to be maximised rather than a property to be designed, and it will produce bloated, sluggish, capital-hungry operations that are no better at surviving the next shock, which will not look like the last one.

The tell is always the same: an absence of discrimination. The fragile organisation could not tell you which of its economies were dangerous. The panicking organisation cannot tell you which of its new buffers are necessary. Both have skipped the only work that matters, which is the patient, specific, unglamorous business of understanding their own chain well enough to know where failure is cheap and where it is ruinous.

7. Pricing the failure

We should stop asking whether to be resilient or efficient. The question assumes they are rivals, and the assumption is the disease. Resilience, properly understood, is efficiency — efficiency measured over the whole life of the system and the whole distribution of its possible futures, rather than over a single quarter's landed cost.

The practical shift is smaller than the rhetoric suggests and harder than it sounds. It is to put a number, however rough, on the cost of failure, and to let that number sit on the same page as the cost of the part. The moment the eleven-week line stoppage appears next to the 0.36 per cent saving, the decision makes itself, and no dial is required. What made us fragile was never efficiency. It was a way of measuring efficiency that had quietly agreed not to look at the most expensive thing that could happen.

The organisations that come out of this era stronger will not be the ones that chose resilience over efficiency, nor the ones that spend the next two years over-correcting. They will be the ones that finally started costing the failures they had trained themselves not to see — and discovered, as they always do, that the cheapest supply chain and the sturdiest one were the same design all along.