

The factory is not there to make things

There is a third answer when a critical category is concentrated and the only options on the table are buying the supplier or living with the dependency. Own a slice of production — sized on what keeps a line alive rather than on what it supplies — and use it to make everything outside it perform better.

The board paper recommends acquiring the supplier. The chief operating officer does not want a factory, and is right not to want one: the seller knows exactly how dependent you are, which is why the price is what it is; integrating a plant will consume leadership attention during the eighteen months you can least afford to lose it; and the thing that made the supplier worth buying — a team and a body of process knowledge — is the part most likely to walk out during the first year of new ownership. What you would be left holding is a building and a headcount number.

So the alternative on the table is to accept the dependency and manage it with contracts. That is also a bad answer, for the reason everybody already knows.

Both options are answers to the same question, and the question is wrong. It asks how much volume you should make yourself. The better question is what a line of your own would teach you — because the answer to that one is not a volume at all.

Own a fifth to a third. Use it to improve the rest.

Own twenty to thirty per cent of a critical category's production capacity, and use it to make the other seventy or eighty per cent perform better. That is not vertical integration and it should not be sold as one. It is a capability acquisition, and the distinction matters commercially, because the two are valued differently and integrated differently.

The output of the line is not the product. It is the standard everything else gets measured against.

A line you own builds to your specification rather than to a supplier's interpretation of it. Every unit tests your interface definition in practice, which is the only place interface problems actually surface. It gives product development somewhere to take a change without joining a supplier's queue. It converts factory knowledge into commissioning competence when the same people go to site. It gives your engineers standing when they walk into a supplier's plant, because they are speaking as people who build the thing rather than as customers with a checklist. And it gives you somewhere to put an emergency — the job a commercial supplier would refuse, price punitively, or accept and miss.

That last one is worth pausing on, because it is the role nobody underwrites and usually the one that pays for the whole line. Its economic case is not unit cost. It is **averted delay cost**, which on a project of any size is an order of magnitude larger and lands in a single quarter.

The cap is the design, not a compromise

Everything above argues for building something. Nothing above argues for building a lot, and the discipline of this model is the deliberate refusal to.

Stay below half and the line never has to run at high utilisation to justify its own existence, so a soft quarter is an operational annoyance rather than a facilities crisis. It stays small enough that manufacturing headcount does not become an industrial-relations exercise during your steepest growth. It does not require the different, heavier discipline of running a majority-volume plant, which is a genuinely different job competing for the same scarce leadership attention. And it leaves your incumbent suppliers with a business worth having — which preserves the resilience you built the line to protect rather than destroying it in the name of securing it.

Most vertical integration is justified by scale. This is justified by refusing scale, and that inversion is the argument to lead with in the room.

Where the floor actually comes from

A fifth to a third is a useful shorthand and a poor target, because a percentage cannot be a design input. The line has to be sized on an absolute volume — the smallest volume at which it is genuinely a line rather than a demonstration — and the percentage is whatever that turns out to be when you divide it by the category's demand.

Three things set that floor, and you take the largest of them.

Skill currency. The interval between builds has to stay inside the period over which your crew remains proficient, which is a property of your product and your process and which you can measure. A line that builds twice a year does not hold a capability; it holds a memory of one, and relearns on every unit.

Crew indivisibility. You cannot employ two-thirds of an integration crew. The minimum team is set by the work content of a single unit — the stations that have to be manned to build one — not by how many units anyone wants. Below the volume that occupies that crew at acceptable utilisation, you are paying for idle capability, and idle capability is the first thing cut in the first bad quarter.

Benchmark validity. The line has to be far enough down its own learning curve that its yield and cycle time are numbers a supplier can be held to. Presenting a first-article result as the standard is not a benchmark, and any competent supplier will say so, at which point the central purpose of the line is gone.

Then add headroom, because a line running at its planned load has no surge role. Capacity has to be sized above planned volume or the fifth function does not exist.

The percentage is an output. If you set it as a target you have started at the wrong end.

And sometimes the arithmetic says no

Divide that minimum volume by the category's demand and one of two things happens.

If the answer sits comfortably below the ceiling, the model works and the band you have heard quoted is simply what the arithmetic produces at that scale. If the answer comes out near or above the ceiling — if the smallest viable line would be forty per cent or more of everything you buy in that category — then the category is too small for this play. You cannot have the benefits without taking on the fixed cost, the distraction and the supplier damage the cap exists to avoid, and pretending otherwise gets you the worst of both.

In that case you are back to the two options you started with, and the honest thing is to say so early rather than to shave the design until it fits.

Write the downturn rule before the downturn

Here is the case the cap does not cover. Hold the line at its absolute volume, halve the category's demand, and your share doubles — straight into the territory the ceiling was set to keep you out of. Nothing has gone wrong operationally. The arithmetic simply moved underneath you.

What happens next is usually not a decision. It is a reflex: protect the team you own and take the volume out of the external base. That is understandable, and it dismantles the supplier resilience the whole model was built to preserve, at the precise moment you need it. The opposite reflex is no better — flex your own line down to protect suppliers, drop below minimum viable, and the capability you paid for is gone, at a rebuild cost that exceeds anything the flex saved.

So the split has to be agreed while everyone is calm: what proportion of any volume reduction the in-house line absorbs, written into the sourcing agreements and into the board paper that approves the investment. It costs something real — you will deliberately underload a line you are paying for — and it buys the thing the model exists for, which is a supply base that is still there on the way back up.

Two instruments, one objective

This sits alongside the other way of keeping control of a product you do not build. Where a category will accept an interface-defined relationship, hold the boundary: own the interface definition and the qualification gate, and let the supplier own what happens inside. Where the market will not accept that and insists on full manufacturing prints, an interface definition will not hold — and owning enough of the line yourself means you never have to hand the prints over.

Same objective, opposite instruments, and which one you reach for depends on what the category will bear. Naming the choice explicitly is worth more than either technique on its own.

When not to do this at all

If the only argument is unit economics — supplier margin looks recapturable at volume — then this is not a capability acquisition and should not be dressed as one. You are buying a factory, the conventional risks apply in full, and they should be priced honestly.

And if you do go ahead, hold on to what you decided you were buying. If the value is the team and the process knowledge rather than the building, the transaction has to say so: what the retention terms are, who the named people are, and what happens if they leave in the first year. A capability acquisition that loses the capability has bought the one asset it explicitly did not want.

That is the work I do.

WHERE THIS GOES NEXT

Where you build it — and how you can defend it →

The manufacturing strategy and footprint method — product, process, make-or-buy, flow, capacity, ownership, geography and organisation, in that order, with make-or-buy decided per step against strategic, capability, economic and risk. Or start a conversation: rtackconsulting.com/contact