

The forecast will be wrong. Do it anyway.

Everyone knows the number is wrong, and most operating teams draw the wrong conclusion from that — they stop defending it, and quietly stop doing the work behind it. The forecast is disposable. The forecasting is not, because it is the only process that makes a company argue about next year before next year happens.

Ask a management team how good their forecast is and you will get a wry answer. Ask what they do about it and you get a shrug: sales submit a number nobody believes, operations build to a different one they do not publish, finance holds a third. Nobody is lying. Each is right about the part of the business they can see.

The damage is not that the forecast is wrong. It was always going to be wrong. The damage is that three plans are running, none has been argued out loud, and the differences surface only when something has already failed — a line stops, a customer is given a date that was never achievable, or material lands against demand that evaporated two months ago.

The plan is disposable. The planning is not.

That distinction is the whole subject. A plan is a snapshot of what you currently believe, and it starts decaying the moment it is published. Planning is the recurring act of making people who see different parts of the business reconcile what they believe, on a date, in front of each other. The output you keep is not the number. It is the argument you had to have to produce it, and the fact that the disagreements are now on a page instead of in three heads.

What the process is actually for

Sales and operations planning has an unhelpful reputation, largely because it is so often installed as a monthly meeting rather than as a process. A meeting produces attendance. A process produces decisions, and it needs the things decisions need: named roles, a fixed calendar, an agreed set of inputs, and somebody with the authority to settle a disagreement in the room.

The rhythm that works is monthly and runs in four passes. **Demand** first — what does the market look like, what has changed, what are we now saying we will sell. **Supply** second — can we build that, with what we have and what we can get, and where does it break. **Finance** third — what does that plan do to revenue, margin, inventory and cash. **Then the decision** — because the first three passes will not agree, and someone has to choose.

The order matters and the separation matters. Run demand and supply in the same session and the supply constraint quietly edits the demand number before anyone has had to defend it. That is how a company ends up forecasting its own capacity back to itself and calling it a market view.

A single number is a wish

The first practical change is to stop asking sales for a number and start asking for three. Demand should carry volume, timing *and* probability — what is sold, what is likely, and what is possible — because each tier buys something different and the difference is what you plan against.

What is sold justifies committing money. What is likely justifies committing long-lead material and holding capacity. What is possible justifies nothing at all except knowing it is there, which is exactly the point: an opportunity nobody has quantified still consumes attention, and putting it in the bottom tier is how it stops distorting the plan while remaining visible.

It is also the answer to the objection that forecasting a bespoke, lumpy project business is meaningless. It is meaningless as a single number. As three tiers with dates against them it is the most useful document in the building, because it separates the safe commitments from the bets — and makes whoever is placing the bet say so.

And your supply plan needs the same discipline

Less obvious, and more often skipped: the supply side has to separate what is **contracted** from what is merely **expected**. A supplier's indication of what they could probably do next year is not the same object as a purchase order or a capacity agreement, and a plan that adds them together is not a plan, it is optimism with a total at the bottom.

Put the two sides on one sheet and the arithmetic does something no discussion can. It produces a single line — what you can actually promise — and it produces it by month rather than as an annual total, which is where the annual total was hiding the problem.

Available to promise — illustrative

One product family, twelve months, units. Demand carries a probability tier. Supply separates what is contracted from what is merely expected.

	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	YEAR
DEMAND													
Sold — firm order	40	55	60	45	20	10	—	—	—	—	—	—	230
High probability	—	5	15	35	55	60	50	40	25	15	5	—	305
Low probability	—	—	—	5	15	30	45	60	70	75	80	85	465
Total demand	40	60	75	85	90	100	95	100	95	90	85	85	1,000
SUPPLY — CONSTRAINED BY THE LONG-LEAD MODULE													
Contracted	45	60	75	70	60	55	40	20	—	—	—	—	425
Expected, not contracted	—	—	—	—	20	35	50	75	100	100	100	100	580
Total availability	45	60	75	70	80	90	90	95	100	100	100	100	1,005
POSITION													
Movement in month	5	—	—	(15)	(10)	(10)	(5)	(5)	5	10	15	15	5
Available to promise / (short)	5	5	5	(10)	(20)	(30)	(35)	(40)	(35)	(25)	(10)	5	5
BINDING CONSTRAINT													
	—	—	—	Module	Module	Module	Module	Module	Module	Module	Module	—	

Illustrative figures. Two things the sheet makes unarguable: the shortfall is visible in January, three months before it bites — and the recovery from September is built entirely on supply nobody has contracted for.

Two things fall out of a sheet like that, and neither is available from any other document. The shortfall in April is visible in January, which is the entire value of the exercise — three months is the difference between requalifying a second source and paying for air freight. And the recovery from September rests entirely on supply nobody has contracted for, which is a sentence a board can act on and a number on its own never says.


Model three or four constraints, not all of them

Modelling every constraint is impossible, and the attempt is where these exercises die. Take the top three or four out of the strategic supply chain plan — and they are rarely all of the same kind. A capacity limit inside your own plant: test bays, a coating line, the commissioning crew. A bottleneck at a supplier, upstream of anything you control. And the open supply commitments themselves, which constrain what you can promise whether or not anything is physically scarce.

Then show them by month, because **the binding one moves**. The constraint that dominates the first quarter is frequently irrelevant by the third, and the one that bites in the third is invisible today unless you are deliberately looking for it. A model that names a single permanent bottleneck is describing last year.

Your ramp against one supplier's capacity — illustrative

One long-lead item. Your projects stacked by how likely they are, against what the supplier says they can build.



Illustrative figures. You do not need to know the supplier's other commitments for this to be actionable — you are over their line on your own demand, four months out.

The supplier case is the one worth drawing, because it is the one nobody owns. Stack your own ramp by probability tier, put the supplier's stated capacity across it as a line, and a breach appears several months out — **on your demand alone**. That qualifier is the whole value. You will never know their other commitments and you do not need to: if you are over their line on your own projects, the conversation has to happen whatever else they are carrying, and it can happen while there is still time to qualify another source, move a project, or tell a customer early rather than late.

They will be wrong. But they will make you think.

Which is the same argument as the forecast, one level down. Nobody sensible expects the constraint model to be right. Its value is that it puts three or four candidates on the table with dates against them, and forces the question everyone avoids: what are we doing about that one, specifically, this quarter. A named constraint is a problem you can assign to somebody. An unnamed one is a surprise with a date on it that nobody has read.

How far out to plan

The horizon is not the financial year. It is a property of your supply base, and the working rule is roughly ten times your longest lead-time item — because you need enough cycles inside the window for a decision made today to be capable of changing anything.

A six-week critical item means planning a year out. A nine-month casting means several years, whether the company admits it or not — and if the window is twelve months, every decision about that item is already made by the time it appears in the plan. Set the horizon by the constraint. If that produces an uncomfortable number, that is the finding.

The events that will hurt your plan are the ones outside your control. Sequence the work outside-in.

Which is the argument for building the process from the outside in. The things most likely to break next year's plan are your customers' decisions, your competitors' moves and your suppliers' constraints — none of which are visible in your own history. Use history where you have it and it is stable. Where you do not, use a forward view and re-cut it often, because a fast wrong answer that gets corrected monthly beats a slow one that is defended annually.

What makes it fail

Three failure modes, and they are consistent enough to check for in the first week.

Nobody owns the output. A planning system produces suggestions, and suggestions with no owner become noise inside a quarter. One named person — a master scheduler by whatever title — has to review what the system proposes and own what is published. Without that role the process degrades into a report nobody reads and everyone is entitled to disagree with.

Planning at the wrong level. Forecast product families, not every part number. Forecasting at item level is where the effort goes to die: it consumes the planners, produces false precision, and delays the plan past the point of usefulness. Get the family right and let the detail be derived.

Waiting for the system. A spreadsheet runs an honest process for longer than people expect — past the first year, often the second. Beyond roughly fifty or sixty planned items with multi-level bills of material and more than two or three significant customers, the arithmetic stops fitting in a spreadsheet anyone can audit. That is the trigger to buy. Buying earlier does not accelerate anything; it makes your current process permanent and expensive. Install the people first, then the tool.

And it is worth more outside the building than in

Suppliers allocate capacity to customers whose demand signal has been consistent enough to plan against. Customers are judging whether your date is real. Both are credibility judgements, formed over quarters. A company that publishes a plan, misses it, explains why and republishes is easier to supply and easier to buy from than one that publishes nothing and then asks for a favour in a shortage.

When this is not the problem

If you have one product and one customer, you do not have a planning problem. You have a concentration problem, and no amount of process makes a single customer less than a single customer.

And if the plan is sound but nothing arrives on time, the constraint is execution or supplier performance, not planning. Re-running the forecast will not fix a supplier who misses. It will only give you a more accurate description of missing.

That is the work I do.

WHERE THIS GOES NEXT

It's hard to make forecasts, especially about the future →

The planning, visibility and risk method — sales and operations planning with roles and a calendar, demand carrying probability, exception reporting on a fixed cadence, and work-in-progress exposure modelled before the event. Or start a conversation: rtackconsulting.com/contact