

DIAGNOSTIC BRIEF · PART TWO OF A FOUR-PAPER SERIES

The Forecast You Can't Trust: How Period-to-Period Volatility Becomes Inventory You Can't Sell and Expedites You Can't Avoid

A perspective for Boards, CFOs, and Senior Supply Chain Leaders



Why is this brief

The first paper in this series argued that the trillion-dollar inventory burden carried by manufacturers and distributors is not a forecasting problem to be solved with better tooling, but an operating-model problem to be solved with a different approach to replenishment. This brief is the next step. Before a board can endorse a meaningful change in how the supply chain is run, it needs to see, in concrete terms, what is actually wrong with the system in front of it. The most direct way to do that is to look at the forecast — not as a number, but as a moving picture.



Erik Bush

Co-founder and CEO of Intuiflow

1. What Forecast Variability Actually Looks Like

The chart below is a forecast waterfall of an item a Tier 1 automotive supplier ships to an automobile assembly plant — a view almost no senior executive ever sees, yet one that makes the problem unmistakable. Each column represents a single future week of demand. Each row shows the forecast that was issued for that week at a different point in time across a thirteen-week planning horizon, ordered from the oldest forecast at the top (made thirteen weeks before the target week) to the most recent at the bottom (made one week before).

Cells shaded red indicate forecasts that fell more than fifteen percent below the column average over the thirteen-week horizon; cells shaded green indicate forecasts that ran more than fifteen percent above. Beneath the vintage rows, the actual demand realized is shown first, separated from the summary statistics that follow: the average forecast over the horizon, the minimum and maximum forecasts issued, and the variance of those extremes against both the horizon average and the realized demand.

The bottom row reports the variability range — the spread between the maximum and minimum variance to actual, expressed in percentage points — and shades it red, amber, or green according to the interpretation thresholds defined in Section 4.

Vintage forecast	Week 4	Wk 5 (Focal)	Week 6	Week 7
L-13 (13 weeks before)	68	65	72	70
L-12	60	78	68	88
L-11	78	60	80	65
L-10	65	72	55	78
L-9	50	85	78	72
L-8	75	71	65	80
L-7	70	88	70	72
L-6	82	87	75	68
L-5	60	54	82	75
L-4	38	45	60	90
L-3	55	42	72	65
L-2	72	45	48	78
L-1	90	100	85	60
Actual demand	55	50	45	70

Vintage forecast	Week 4	Wk 5 (Focal)	Week 6	Week 7
Average over horizon	66.4	66.8	70	73.9
Minimum forecast	38	42	48	60
Maximum forecast	90	100	85	90
Min variance to average	-43%	-39%	-31%	-19%
Max variance to average	+36%	+46%	+21%	+22%
Min variance to actual	-31%	-16%	+7%	-14%
Max variance to actual	+64%	+100%	+89%	+29%
Variability range (Max – Min vs actual)	95 pts	116 pts	82 pts	43 pts

Illustrative forecast waterfall – automotive Tier 1 seat frame supplied to OEM assembly plant. Units in thousands; variance and variability rows in percentage points. Week 5 figures (L-8 through L-1) drawn from a real client case; other columns shown for pattern comparison.

Read the Week 5 column from top to bottom. Thirteen weeks before the part was needed, the forecast called for 65,000 units. It drifted through 78,000, 60,000, 72,000, and 85,000 across the next five vintages before settling at 71,000 with eight weeks to go. From there it climbed to 88,000 and 87,000, then fell hard – to 54,000, 45,000, and 42,000. With one week to go, the forecast suddenly spiked to 100,000.

The actual demand realized in Week 5 was 50,000 – close to the minimum of the range, half of the most recent forecast, and roughly twenty-five percent below the horizon average. The variance statistics make the dispersion concrete: against the realized demand of 50,000, the minimum forecast was off by 16 percent and the maximum by a full 100 percent – a variability range of 116 percentage points.

Now look at Weeks 4, 6, and 7. The amplitudes differ, but the pattern does not. Week 6 even shows a chronic over-forecast bias – its minimum forecast still landed seven percent above the actual. Three of the four weeks fall in the structural zone for variability; the four-week average sits at 84 percentage points.



Everything is being tried, they're trying to coordinate it, but they're changing their mind because of the noise in the forecast."


2. Why the Variation Is Structural, Not Fixable

In a Tier 1 supplier, this forecasted demand is coming directly from the automotive assembly plant using this part. It represents their changing requirements as they strive to address imbalances in the inventory of completed vehicles relative to evolving market demand.

In other settings, the same volatility appears in the evolving SKU-level forecasts a company generates for the end items it sells to the market. The natural reaction to a chart like this is to ask which planner, which algorithm, or which data feed is responsible for the volatility. It pressurizes the quest to make the forecast more accurate — and potentially, to assume the next solution investment will smooth it out. Thirty years of forecasting investment across U.S. manufacturing have shown that this is not the case. The volatility is structural. Four mechanisms keep it that way.

Aggregation hides it. A category-level or division-level forecast can appear reasonably stable while the SKU-level forecasts inside it vary week to week, because the smoothing happens at the rolled-up level, not at the level where supply decisions are actually made. Many planning organizations report forecast accuracy at the aggregation level that makes the metric look good without addressing the inaccuracies at the item level — which is where the cost lives.

Promotion and program timing creates it. Marketing campaigns, channel commitments, and customer-driven orders arrive late, change late, and are never as well-coordinated with supply as the official process implies. Each late signal triggers a forecast revision.



Demand sensing amplifies more than it dampens. Modern demand-sensing techniques react quickly to recent signal, which is intended to make forecasts more responsive but in practice converts short-term noise into permanent record. The system catches the swing and writes it into the plan.

Human override usually makes it worse. Planners, whose job is to make the forecast more accurate, either feel compelled to make adjustments or do so because they do not trust the forecast the system is generating.

While the automotive waterfall above is one industry, the same patterns appear in fashion apparel, consumer electronics, food and beverage, and industrial distribution. Different industries, same physics.

3. How Forecast Variability Transfers Directly Into Supply Requirements

05

The reason forecast variability matters for the board is that it does not stay in the planning system. It propagates, with very little dampening, into every downstream supply decision.

Each new forecast triggers a new MRP run. Each new MRP run generates a new production plan. Each new production plan revises supplier releases. Each revised release either commits the supplier to product they may not need or pulls forward product that was already in flight. The signal that started as a number on a planner's screen ends up as a change order on a supplier's shop floor, a different shift pattern at the assembly plant, or an expedited freight shipment booked at premium cost.

The financial consequences compound in three directions at once. Inventory ratchets up on items with actual demand well below forecast, while concurrently items that are exceeding forecast quietly become stock-out exposures. Expedites multiply because every upward forecast revision close to the demand week triggers a scramble to catch up — air freight, weekend overtime, premium supplier surcharges. And capacity gets whipsawed: weeks of overtime followed by weeks of idle time, both expensive, both eroding workforce stability.

The most damaging consequences are significant: the use of production resources to make items not selling to the forecast; the expedited freight spend to bring in materials not ordered on time; the planner workload expended trying to control the bullwhip effect in planning. In the end, order fill rates don't improve and inventory turns remain stagnant.

4. The Diagnostic: Build Your Own Waterfall

The exercise below can be run by any supply chain organization that retains its historical forecasts. It takes a competent analyst less than a day for a single product family, and it produces a result that no executive presentation can refute, because the data comes from the company's own systems.

- ✓ Select fifty to one hundred SKU-location combinations that represent the company's highest-revenue or highest-inventory exposure. Pareto analysis usually identifies these quickly.
- ✓ For each combination, retrieve every weekly or monthly forecast issued over the thirteen-week horizon, aligned to the same target period.
- ✓ Lay each combination out as a waterfall matrix in the form shown above, with the actual demand at the top of the analysis block followed by the average, minimum, and maximum forecasts and their variances against both the average and the actual.
- ✓ Compute the variability range for each target period: the maximum variance from actual demand minus the minimum variance from actual demand, expressed in percentage points. In the example above, Week 5's maximum forecast ran 100 percent above the realized demand while its minimum forecast came in 16 percent below — a variability range of 116 percentage points. Repeat across target periods and SKUs to produce a single comparable number for each.

Reporting it to the board


Report two summary measures across the SKUs analyzed. The average variability range gives the systemic picture — what the typical forecast costs the supply chain. The worst-case variability range — the SKU or week with the widest spread — gives the boardroom soundbite that resists dismissal: “for one in five of our top SKUs, the forecast swung by more than one hundred percentage points before reality arrived.” Senior leaders accept averages with a shrug; they cannot ignore tail cases described in their own data.

Interpretation thresholds

The interpretation is straightforward. Below thirty percentage points of variability range, the forecast is reasonably stable and the operating model is probably not the bottleneck. Between thirty and seventy-five points, the forecast is doing meaningful damage downstream, even if the headline accuracy metric looks acceptable. Above seventy-five points, forecast-led planning is structurally misaligned with the demand pattern of the business, and continuing to invest against it is compounding the working-capital and expedite costs documented above. In the example waterfall, three of the four target weeks fall in the structural zone, and the four-week average — 84 percentage points — confirms that the supplier is operating well past the point where additional forecasting investment can solve the problem.

5. What This Sets Up

The waterfall is meant to settle an argument, not to start a new one. If the variation in the forecast is structural — and the four mechanisms above ensure that it is — then no amount of additional investment in forecasting accuracy will eliminate it. The only durable response is to design an operating model that absorbs the variation rather than transmitting it. Where the buffers sit, how they are sized, what triggers replenishment, and how forward scenarios are tested against the model become the questions that matter. That is the subject of the third paper in this series.



About this brief

This is the second paper in a four-part series on supply chain operating model design. The first paper made the strategic case for shifting from forecast-led to demand-driven planning. This brief equips senior leaders to diagnose the problem in their own data. The third paper covers the design of the demand-driven operating model; the fourth addresses deployment and the evolution of the planning function.