POSITION PAPER

The Trillion-Dollar Forecast: Why Better Demand Planning Won't Free Your Working Capital — But Better Operating Models Will

A perspective for Boards, CFOs, and Audit & Risk Committees

algo

Executive Summary

U.S. manufacturers are sitting on roughly one trillion dollars of inventory — more, relative to sales, than they were thirty years ago. Three decades of investment in forecasting software, demand planning teams, S&OP rituals, and ERP modernization have not bent the curve. For boards and CFOs, this is not a supply chain problem. It is a capital allocation problem.

This paper argues that the persistent inventory burden carried by manufacturers and distributors is the predictable consequence of a planning paradigm built on two assumptions that no longer hold: that forecasts can be made accurate enough at the SKU level to drive replenishment, and that the supply chain operates without meaningful variation. Both assumptions are demonstrably false, and continuing to invest against them produces diminishing — and eventually negative — returns.

The path to working capital release runs not through more sophisticated forecasting, but through redesigning the company's operating model so that replenishment responds to actual demand rather than to a forecast about it. Forecasting is not abandoned in this model — its role is repositioned.

Rather than driving day-to-day replenishment, the forecast becomes the instrument by which leaders test the operating model itself: simulating forward to surface stock-out exposures, pressure-testing rough-cut capacity, and prompting deliberate configuration changes in anticipation of future events, all while the business continues to replenish against actual consumption. Freeing even ten percent of the inventory currently on U.S. manufacturers' balance sheets would unlock approximately one hundred billion dollars of working capital. That is the prize.

This paper outlines how boards and CFOs should think about it, and the questions they should be asking their operating teams.

Erik Bush
Co-founder and CEO of Intuiflow

1. The Forecast Accuracy Trap

The dominant logic in supply chain planning today — embedded in virtually every enterprise system on the market — is that the right level of inventory can be derived from a sufficiently good forecast. Plan against the forecast, schedule backwards from lead times, and the stock will be there when the customer wants it.

The premise is intuitive. It is also wrong in two specific and expensive ways. First, demand forecasts are not accurate at the level customers actually buy. Consumers do not buy categories from Coca-Cola; they buy the specific product, in the specific size, in the specific channel. The further a forecast is decomposed toward that buying unit, the more error it encounters. This is not a tooling problem solvable by adding data, machine learning layers, or planner headcount. It is a property of the system.

Second, classical planning logic assumes the world runs on time. Lead times are achieved. Promotions are scheduled with adequate time to respond. Suppliers ship on the required date. These factors, and others are not consistently reliable, and the response — layering safety stock on top of forecast — produces what practitioners call a bimodal inventory distribution: too much of the wrong stuff and too little of the right stuff, simultaneously, on the same balance sheet.

The lived consequence is a treadmill familiar to every operating leader: weekly forecast revisions that swing thirty to fifty percent against a customer's own confirmed orders; planners burning out and turning over because the job is, in their words, living on a razor's edge; spreadsheets quietly running the business underneath the official system; and capital expenditure programs — plants, lines, distribution centers — sized against numbers no one fully trusts.

I don't know how anybody in a forecast-driven business knows when to stop working on their demand plan. You can always make it better — but at what point does the incremental effort become counterproductive?"

The honest answer, in most companies, is that they never stop. Forecast accuracy follows a steep curve of diminishing returns: each additional point of accuracy costs disproportionately more than the last — in data, in tooling, in headcount, in management attention — while the inventory benefit shrinks. Several large enterprises have recently concluded the math no longer works. A major global consumer products company recently shut down a long-standing centralized demand planning organization after determining the spend was not producing measurable improvement

Others continue to staff up, hiring "armies of planners" to scrutinize SKU combinations whose run rates do not justify the attention. Boards rarely see this trade-off articulated in capital terms. They should.

2. A Different First Principle: Plan the Model, Not the Forecast

The alternative is not to abandon forecasting. It is to reposition it. The question to ask is not "how accurate is our forecast?" but "how do we ensure stock is available to serve actual demand, with the smallest working capital footprint possible?"

Demand-driven approaches — rooted in the same Kanban logic Toyota proved out half a century ago — answer that question by inverting the flow. Replenishment is triggered by what customers actually consumed, not by what a forecast predicted they would. Strategic inventory buffers, sized for each item's real demand rate and variability, absorb the volatility that classical planning tries unsuccessfully to forecast away. Day-to-day execution is governed by real signals — orders, shipments, point-of-sale — not by a number that was written down weeks earlier and is already drifting.

The business model that already operates this way is more familiar than most boards realize. Zara, the apparel company most often studied as a counterexample to forecast-led production, deliberately seeds the market with limited quantities, observes what sells, and then replenishes rapidly against real consumption. Its competitors bet the season on forecasts and write down the unsold inventory. The financial gap between the two approaches — in inventory turns, in margin retention, in capital intensity — is now well documented. The principle generalizes well beyond apparel.

The Repositioned Role of the Forecast

If replenishment no longer waits on the forecast, what is the forecast for? Its role becomes evaluative rather than directive. It is the instrument by which leaders test whether the operating model they have configured is fit to absorb what the business expects to face. In practice, this looks like an S&OP discipline reoriented around three questions.

Will our current inventory buffer configuration hold stock availability under the demand we expect over the planning horizon? Will the projected workload — production hours, line capacity, supplier draw — fit within our rough-cut capacity, or are we walking into a constraint? And where are the specific items, locations, or windows in which our model, as it stands today, is likely to fail?

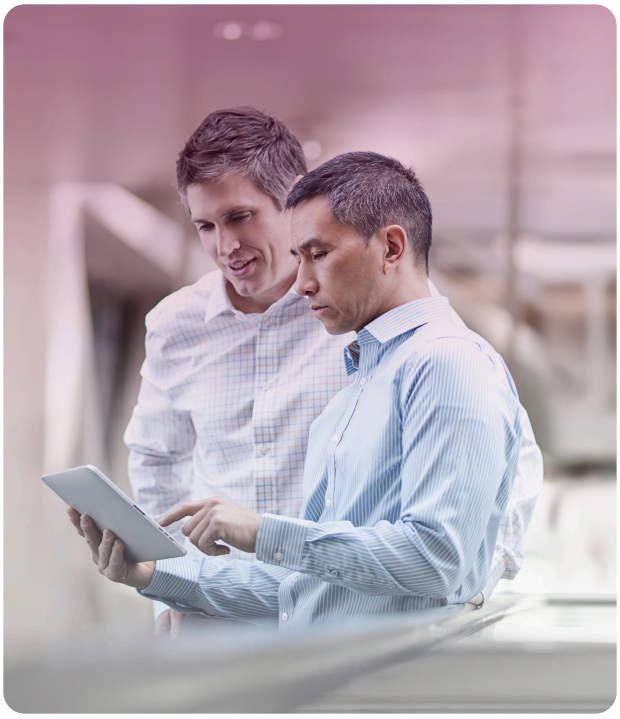

Modern demand-driven systems answer these questions by ingesting one or more forecasts — typically a base case alongside high and low scenarios — and simulating their effect on the operating model day by day. Buffers flex. Projected inventory positions are surfaced. Stock-out exposures appear with names, dates, and quantities attached. Capacity loads are projected forward so production and supply teams can see where the model will strain. The output is not a production schedule. It is a diagnosis.

That diagnosis then drives configuration changes — not commitments to build to the forecast, but deliberate adjustments to how the model itself is set up. An item showing exposure under a seasonal ramp might be moved from a trailing rate of demand to a forward-looking one. A buffer profile might be re- calibrated to absorb higher variability. A lead-time assumption might be tightened or relaxed. A new product introduction might be staged with a shorter initial demand window and re-tuned as real consumption arrives. The system continues to replenish based on actual demand. The forecast simply ensures the model is set up to handle what is coming.

> Roughly right is better than precisely wrong."

This is the principle that lets a company stop chasing forecast accuracy. Once the operating model is configured to absorb normal variation and stress-tested against forward scenarios, additional forecast precision yields little incremental inventory benefit. Roughly right, fed into a well-designed model, produces better service and lower working capital than precisely wrong fed into a rigid plan.

For manufacturers and distributors of every size, this reframing changes which investments matter. The leverage is not in the forecasting engine. It is in the design of the operating model: where buffers sit, how they are sized, what triggers replenishment, how forward scenarios are tested against it, and how the organization responds when reality drifts.

3. What Boards and CFOs Should Be Asking

Most demand planning conversations never reach the board. They are treated as operational hygiene, line items in a supply chain budget, or vendor selection decisions delegated to IT. That framing understates the stakes. A company carrying twenty percent more inventory than it needs is funding that excess with working capital that could service growth, return to shareholders, or absorb a downturn. Five questions sharpen the conversation:

- What is the marginal return on our next dollar invested in forecast accuracy? If management cannot answer this in working capital or service-level terms, the program is operating without a stopping rule — and is almost certainly past the point of diminishing returns.
- Where on the spectrum from make-to-order to make-to-stock does our business actually sit, and is our planning approach matched to it? A mismatch — forecast-led planning in a high-velocity, high-variability environment — is the most common and most expensive misalignment in supply chain operations.
- How much of our planning effort is concentrated on SKUs that drive a meaningful share of revenue? Pareto analysis routinely shows that a small percent of SKUs generate eighty to ninety percent of revenue, yet planning attention is spread evenly across the long tail. That is wasted capacity.
- Are we trapped in customization? Heavily customized enterprise platforms — the kind sold on the promise of "we can configure it to do anything you want" — narrow the field of vision to the company's own past practices and lock out external innovation. The customization spend is rarely visible at the board level; the strategic cost is significant.
- If we released ten percent of our inventory, what would we do with the capital? If the answer is unclear, the inventory question is not yet a board priority. If the answer is concrete — capacity expansion, debt reduction, M&A capacity, dividend support — the demand planning conversation belongs on the agenda.

4. The Capital Allocation Lens

The temptation, when supply chain performance disappoints, is to buy a better forecasting system, hire more planners, or commission another transformation. Each of those moves is defensible in isolation. None of them, on the evidence of the last thirty years, will materially change the inventory-to-sales ratio of a typical manufacturer or distributor.

What will change it is a deliberate decision — made at the level where capital allocation decisions are made — to shift the operating model from forecast-led to demand-driven, accept that forecasts will remain imperfect, and invest the savings in resilience and growth rather than in chasing accuracy that the data will never deliver. The companies that make this shift early will compound the advantage. Those that do not will continue to fund a trillion-dollar industry-wide working capital drag, one quarter at a time.

> The prize is not a better forecast. The prize is the capital that better forecasts have failed to release.

The Capital Allocation Lens

This position paper synthesizes themes developed during a supply chain leadership working session on demand planning strategy. It is intended as a discussion document for board-level conversations on working capital, supply chain operating model design, and capital allocation in manufacturing and distribution.