

Free the Working Capital Your Forecasts Are Tying Up

Synthefy is a foundation model for supply chain demand forecasting, built on a time-series platform deployed at Samsung, NetApp, Deutsche Telekom, and the US Army. See more at synthefy.com.

01 The Problem

Three failures show up in every demand forecast we audit:

- 01 **The model can't see what moves your demand.** Blind to tariffs, commodity prices, and demand shocks, it stays wrong on the 30 to 50% of your catalog that's lumpy or seasonal.
- 02 **No insight into why the forecast moved.** Planners are told the number changed, not the reason they need to make a better call and defend it.
- 03 **Overrides change the number, never the reasoning.** Signal and bias get patched in by hand and lost, so the same miss compounds.


The result is **about 5%** of your lumpy-SKU inventory sitting idle as excess safety stock, which for a \$300M company is roughly **\$1M of working capital tied up** and **\$250K a year** in carrying cost.

02 Why Synthefy

Three things Synthefy does that most models can't.

Accuracy on the SKUs your model gets wrong.


SKU-specific external signals lift accuracy on the half of your catalog that was stuck.



VALUE ~5% of lumpy-SKU inventory released, about \$1M one-time for a \$300M company.

Forecast trust your team can defend.


Every forecast ships with a signal-by-signal breakdown plus a confidence score.



VALUE The carrying cost you stop paying on the inventory you free, about \$250K a year for a \$300M company.

A model that learns from your whole team.

Each override teaches the next forecast. The loop compounds.




VALUE The model gets smarter each cycle: accuracy climbs, bias drops, and manual adjustments fall.

METHODOLOGY Illustrative for a mid-market company at \$300M revenue. Scales with your inventory load and team size. McKinsey, APQC, ISM, BLS.

03

How Synthefy Compares

Synthefy works on top of the tools you already run. Here is what it adds.



CAPABILITY	SYNTHEFY	BLUE YONDER	KINAXIS	STOCKIQ	EXCEL
Time to first forecast	Hours, API-native	12 to 24 months	Months	Weeks to months	DIY
Finds which connected signals causally drive each SKU	●	◐	◐	○	○
Finds cross-product correlations across the whole catalog automatically	●	◐	◐	○	○
Per-forecast numerical signal attribution with a confidence score	●	◐	◐	◐	○
Turns each planner override into a training signal	●	○	○	○	○
COST	\$\$	\$\$\$\$	\$\$\$\$	\$\$\$	\$\$

● BUILT-IN

◐ LIMITED

○ NOT AVAILABLE

Synthefy works from the signals you connect, and only where external forces actually move demand.

We don't promise perfect accuracy. We deliver a defensible forecast that significantly outperforms your incumbent stack, with confidence scores planners can act on.

04 Customer Case Study

"Huge under-forecasting on lumpy items. The factory is always lagging behind."

HEAD OF SUPPLY CHAIN · GLOBAL INDUSTRIAL OEM

- THE CUSTOMER** A multi-billion dollar industrial OEM with a global supply chain across **six product lines and 160 countries**.
- THE CHALLENGE** **200,000+ active SKUs and over 1M product-country combinations**. Half the catalog lumpy and intermittent, stuck at **40 to 45% accuracy** on the incumbent stack.
- THE APPROACH** Phase zero deployment. Added external signals (copper, commodity ETFs, country macros, news, currency, cross-product), zero-shot.

THE OUTCOMES	METRIC	BASELINE	WITH SYNTHEFY
	Accuracy, lumpy SKUs	40 to 45%	~60%
	Time to live	Months of training and tuning	3 days, zero-shot
	Working capital release	Baseline	~\$145M one-time
	Three-year FCF impact	Baseline	~\$255M cumulative

METHODOLOGY McKinsey 5% accuracy-to-inventory bridge applied to roughly half of customer's \$5.8B year-end inventory (lumpy SKU portion). APQC and ISM carrying cost benchmarks.

OUR MODEL UNIQUELY FOUND SOME OF THESE CORRELATIONS

A copper-price signal cut overprediction nearly 6×.

The model found a strong causal link between a key SKU and copper prices — an external input the incumbent never used. Feeding in that single signal collapsed overprediction from **+11% to +1.8%**.

Cross-product learning caught Turkey's October spike.

By borrowing demand patterns from similar products, the model detected geo-specific October seasonality the incumbent missed. It predicted **3.6K** against **13.8K** actual; Synthefy called **12.3K**.

WHAT'S NEXT

Bring us your hardest forecast.

[Book a demo](#) →