Stop Using Averages for Safety Stock

Autonomy Platform · Public Blog

CONFIDENTIAL

Technical March 2026

Stop Using Averages for Safety Stock

A supplier with a 10-day average lead time that varies from 5 to 25 days is not the same as one that consistently delivers in 9-11 days. Distribution fitting changes everything.

The Lie of Averages

Consider two suppliers, both with a **10-day average lead time**:

Supplier A: Consistent

9-11
days delivery range

Normal distribution, σ = 0.5 days

Minimal buffer needed

Supplier B: Volatile

5-25
days delivery range

Lognormal distribution, heavy right tail

Large buffer required

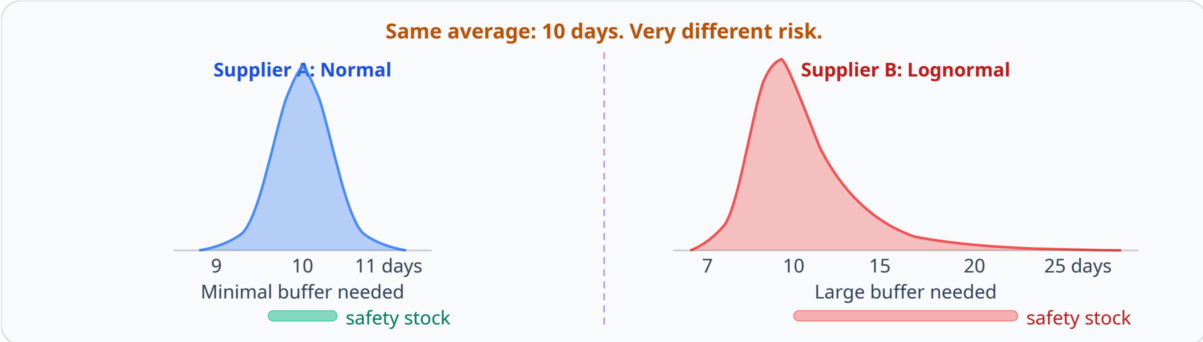

A traditional planning system treats them identically: average lead time, 10 days. But the safety stock required for 95% service level is dramatically different.

Using the average for both means you're either **over-investing in buffer for Supplier A** or **under-protecting against Supplier B**. Both waste money.

The Problem with "Safety Stock"

In traditional MRP, safety stock is treated as a **hard demand target**:

The MRP Safety Stock Trap

1. System generates planned orders to maintain safety stock level
2. These planned orders **compete with real customer demand** for upstream capacity
3. When capacity is constrained, you choose between serving customers and maintaining safety stock
4. **This defeats the entire purpose of having safety stock**

Inventory Buffers: A Different Concept

What we actually need is an **uncertainty absorber**: a buffer that provides protection against demand and supply variability without generating hard demand signals that distort the planning system.

Safety Stock (Traditional)	Inventory Buffers (Autonomy)
Hard demand target → generates planned orders → competes with real demand → distorts system	Uncertainty absorber → soft netting → replenishes when capacity allows → protects real demand

The key difference: buffer-replenishment planned orders get **lower priority** than demand-driven orders. Soft-buffer netting means the system replenishes buffers when capacity is available, not at the expense of real demand. Learn more about inventory buffers.

Distribution Fitting: The Technical Foundation

Autonomy fits actual distributions to operational variables automatically. Instead of assuming one distribution shape fits all, the system identifies the **best-fit distribution** for each variable from your historical data.

Why Distribution Shape Matters

If lead times follow a **Weibull distribution** rather than a **Normal distribution**, the buffer calculation changes significantly. Using the wrong distribution leads to either excess inventory or unacceptable stockout risk.

Normal	Lognormal	Weibull
Symmetric, thin tails	Right-skewed, heavy tail	Flexible shape parameter

This changes **every downstream calculation**: safety stock, reorder points, ATP availability, and capacity requirements all become more accurate when they use the right distribution instead of a Normal approximation. See stochastic planning for the full technical treatment.

Dynamic Buffer Management

Rather than a fixed safety stock quantity, the Inventory Buffer agent continuously adjusts buffer parameters based on:

Actual Demand Variability

Not assumptions: real observed patterns from your data

Realized Lead Time Distributions

Fitted, not averaged: capturing the true shape of variability

Service Level Targets by Segment

Different customer segments get different service commitments

Current Inventory Position

Relative to the demand pipeline, not a static target

From "What's the Plan?" to "What's the Probability?"

Monte Carlo Simulation

Monte Carlo simulation propagates uncertainty through the entire planning engine, generating the scenario data on which agents train and the calibration sets that power conformal prediction.

Conformal Prediction: Distribution-Free Guarantees

Conformal prediction wraps the simulation output in **distribution-free coverage guarantees**. Instead of raw percentiles, every statement carries a mathematical guarantee:

P50 cost: $2.4M, P90 cost: $2.8M
with guaranteed 90% coverage

85% probability OTIF exceeds 95%
a calibrated bound, not an estimate

Service level risk: 12% chance of dropping below 90% in Q3
holds regardless of distribution

This transforms the planning conversation. Instead of debating whether the plan is “right,” you discuss whether the **guaranteed probability distribution of outcomes** is acceptable, and every agent’s likelihood score is trustworthy by construction.

The Bottom Line

Stop doing this:

• Using averages for lead time calculations
• Treating safety stock as a hard demand target
• Assuming Normal distributions everywhere
• Debating whether the plan is "right"

Start doing this:

• Fitting actual distributions to your data
• Using inventory buffers as uncertainty absorbers
• Applying Monte Carlo simulation for scenarios
• Asking "what's the probability the plan works?"

Related reading: the AIIO operating model · the shared world model · the Decision Stream.

See Autonomy in action

Walk through how Autonomy models, executes, monitors, and governs supply chain decisions with autonomous AI agents.

See It Live