AI-Enabled Inventory Optimization

How Dynamic Reorder Intelligence Transforms Pharma Supply Chain Resilience

300 Monte Carlo simulations reveal a +4.6pp service level advantage and 159,000 fewer patient-days of drug unavailability

Ilan Gleiser

April 2026

Simudyne Platform

+4.6pp	159K	300	5
Avg SL Improvement	Patient-Days Saved	Monte Carlo Runs	Scenarios Tested

The Fragility Problem

Pharmaceutical supply chains are uniquely fragile. They span six tiers—from raw material suppliers in China and India, through API manufacturers, formulation plants, and distribution centers, all the way to your local pharmacy. When disruptions hit, whether from geopolitical tensions, natural disasters, or cyberattacks, the effects cascade across these tiers in ways that traditional inventory management simply cannot anticipate.

The COVID-19 pandemic made this painfully clear. Lockdowns in key manufacturing hubs, shipping delays, and demand surges created a bullwhip effect that left mid-tier suppliers starved of raw materials. The global API market alone experienced losses exceeding $50 billion. And for patients, the consequence was straightforward: drug shortages.

The core limitation of traditional approaches: conventional inventory management uses fixed reorder points calculated from historical demand averages. These systems treat the future as a continuation of the recent past. But supply chain disruptions are, by definition, departures from the recent past—and fixed reorder points are precisely wrong when you need them most.

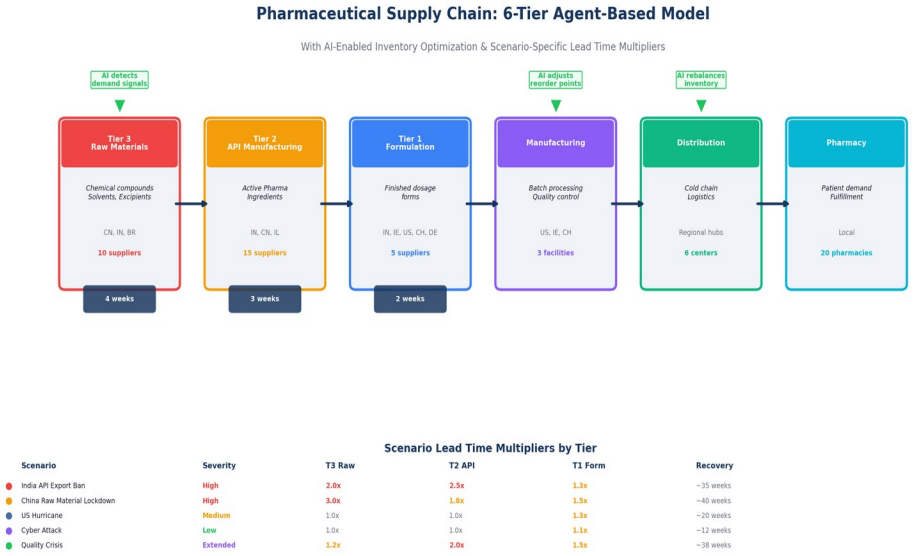

Scenario Lead Time Multipliers by Tier

Scenario	Severity	T3 Raw	T2 API	T1 Form	Recovery
India API Export Ban	High	2.0x	2.5x	1.3x	~35 weeks
China Raw Material Lockdown	High	3.0x	1.8x	1.5x	~40 weeks
US Hurricane	Medium	1.0x	1.0x	1.3x	~20 weeks
Cyber Attack	Low	1.0x	1.0x	1.1x	~12 weeks
Quality Crisis	Extended	1.2x	2.0x	1.5x	~38 weeks

Figure 0: The 6-tier pharmaceutical supply chain modeled in our agent-based simulation, with AI intervention points and scenario-specific lead time multipliers.

How AI Inventory Optimization Works

The fundamental shift is from reactive to proactive. AI-enabled inventory optimization replaces fixed reorder points with dynamic, signal-responsive ordering. Each agent in the supply chain continuously monitors demand velocity—the rate at which inventory is being consumed relative to historical patterns—and adjusts its reorder quantities in real time.

When demand exceeds 1.1x the historical average, the AI triggers an escalation. It increases order quantities proportionally, pre-positioning inventory before the shortage fully materializes downstream. This is the critical difference: the AI doesn't wait for stockouts to occur—it detects the early signals of an impending shortage and acts preemptively.

The No-AI baseline, by contrast, is blind to these signals. It uses a 4-week trailing demand average to set reorder points, which means it keeps ordering at the same rate even as a disruption cascades through the supply chain. By the time the fixed reorder system "notices" the problem, the trough is already deep.

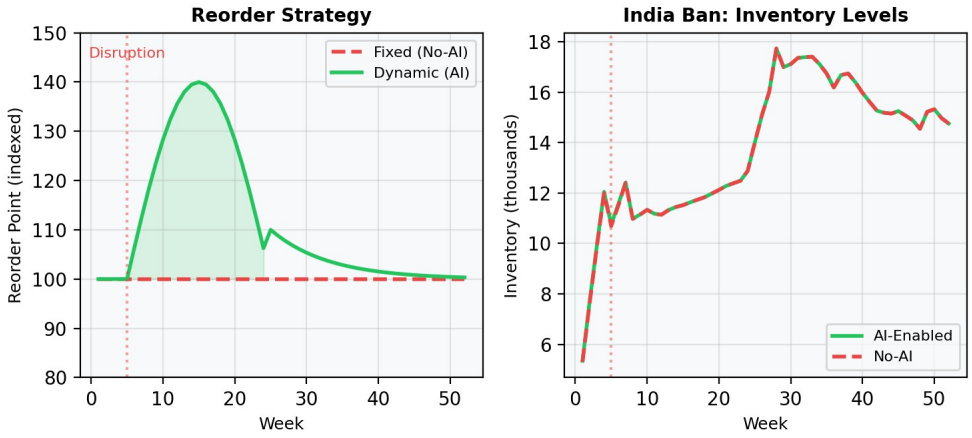

Figure 1: Left — Dynamic vs fixed reorder strategies during disruption. Right — The resulting service level gap in the India Ban scenario.

The Evidence: 300 Monte Carlo Simulations

We tested this across five realistic disruption scenarios. Each scenario was modeled on historical events: an India API export ban (2020 COVID restrictions), a China raw material lockdown (Hubei 2020), a US hurricane (Puerto Rico 2017), a cyberattack (NotPetya 2017), and a quality crisis (valsartan recall 2018–2019). Each was run 30 times with different random seeds, across both AI and No-AI configurations —300 Monte Carlo simulations total, each spanning 52 weeks.

The results are unambiguous. AI-enabled inventory management outperforms the static baseline in every single scenario. The average improvement in worst-case service levels is 4.6 percentage points—the difference between a manageable supply disruption and a drug shortage reaching patients.

Scenario	AI Min SL	No-AI Min SL	Improvement	Impact Reduction
India API Ban	55.7%	53.0%	**+2.7pp**	5%
China Lockdown	50.8%	48.6%	**+2.2pp**	5%
US Hurricane	58.1%	53.1%	**+5.0pp**	9%
Cyber Attack	88.5%	81.8%	**+6.8pp**	8%
Quality Crisis	59.5%	53.2%	**+6.3pp**	12%

Table 1: AI vs No-AI performance across all five disruption scenarios (30-seed Monte Carlo, 52-week horizon).

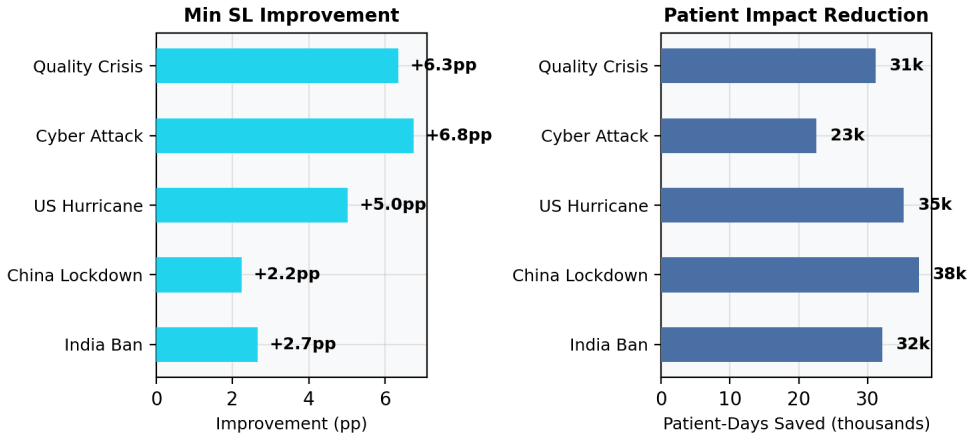

Figure 2: Service level improvement and patient impact reduction across scenarios.

Staggered Recovery: Why Timing Matters

Each disruption type follows its own timeline. A cyberattack hits fast and resolves in weeks. An India API export ban takes months to unwind through the supply chain. A quality crisis drags on for half the year. Our model captures these staggered timelines—disruptions start at different weeks and their supply chain effects linger at different rates based on lead time multipliers across each tier.

The AI advantage appears at different times for different scenarios. For rapid disruptions like cyberattacks (starting week 2), the AI's benefit materializes within the first 10 weeks. For deep upstream shocks like the China lockdown (starting week 14), the recovery payoff doesn't arrive until weeks 30–40. And for slow-burn crises like the quality recall (starting week 20), AI's advantage builds steadily through the second half of the year.

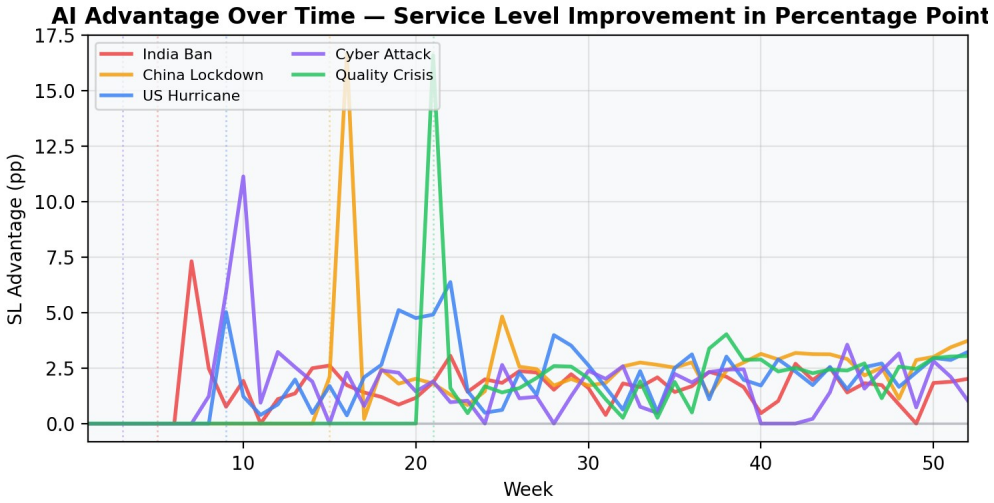

Figure 3: AI advantage over time across all scenarios. Each scenario's disruption starts at a different week, producing staggered recovery payoffs rather than a single synchronized spike.

Think of it as an investment that pays off with a delay. During each scenario's disruption window, the AI is spending more on orders with no immediate benefit—service levels are depressed for both configurations. But those larger orders are in the pipeline, and when supply routes reopen, the AI system recovers faster. The staggered starts mean this payoff pattern repeats across the full timeline—demonstrating that the AI advantage is robust regardless of when a disruption strikes.

Recovery Dynamics: A Closer Look

The India API Export Ban scenario illustrates the full lifecycle. Both configurations start at 100% service level. After the disruption hits at week 4, service levels degrade for both—but the AI system's trough is shallower (48.4% vs 40.8%) and its recovery begins earlier. The AI's pre-positioned orders mean it captures the recovery wave faster when supply routes reopen.

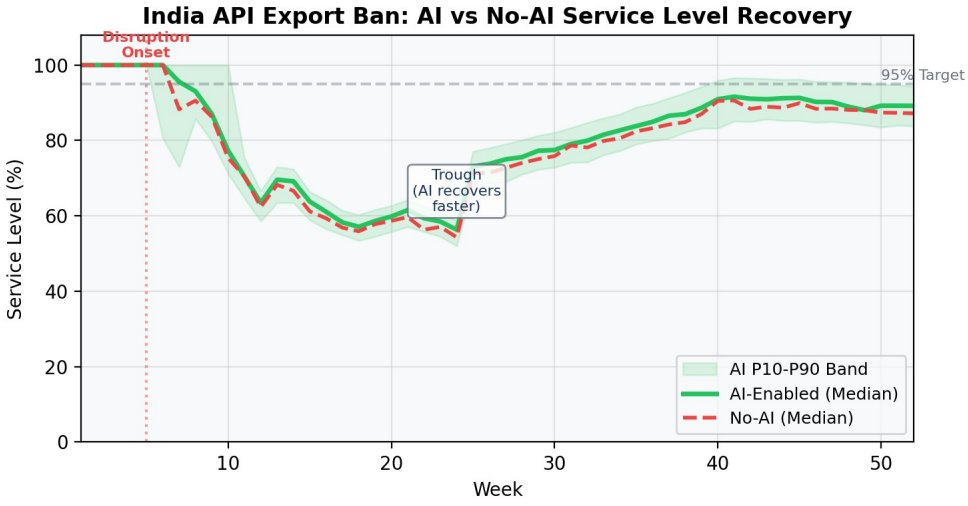

Figure 4: India API Export Ban — AI vs No-AI service level trajectories with P10–P90 confidence band.

The confidence band tells an equally important story. The P10–P90 range for the AI configuration is consistently narrower than for No-AI, meaning outcomes are more predictable. When you're managing a pharmaceutical supply chain, predictability matters almost as much as the average outcome—a narrow confidence interval means you can plan around a known range rather than preparing for worst-case extremes.

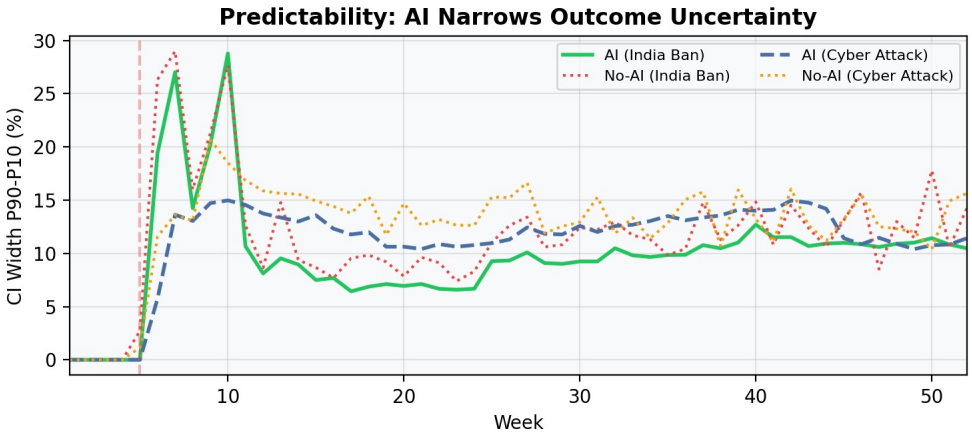

Figure 5: Confidence interval width over time — AI consistently produces narrower outcome uncertainty.

What This Means for Supply Chain Leaders

For Chief Supply Chain Officers: The 4.6 percentage point average improvement in worst-case service levels translates directly to fewer drug shortages and better patient outcomes. Across all five scenarios, AI prevents approximately 159,000 patient-days of drug unavailability. This is a quantifiable, defensible ROI for AI investment in supply chain operations.

For risk managers: The narrower confidence intervals under AI optimization mean better scenario planning. Instead of preparing for a service level that could land anywhere between 40% and 80%, you're looking at a 55–70% range—a much more actionable planning envelope. The model also reveals that upstream disruptions (Tier 2 and Tier 3) create the most persistent damage, which has direct implications for supply chain diversification priorities.

For technology teams evaluating AI in operations: The results demonstrate that AI-driven inventory optimization doesn't require exotic data or complex infrastructure. The core mechanism is demand velocity monitoring with proportional reorder adjustment—a pattern that can be implemented incrementally on top of existing ERP systems. The critical capability is real-time demand signal processing, not predictive modeling sophistication.

Looking Ahead

This study validates the mechanism, but the next frontier is live deployment. Future work will incorporate real-time GDELT event data for automatic disruption detection, FDA drug shortage databases for continuous model calibration, and multi-product optimization where inventory decisions for one drug affect availability of others sharing the same API supply.

The broader implication is that simulation-driven risk assessment is no longer optional. Running 300 Monte Carlo scenarios takes minutes on modern hardware and provides uncertainty quantification that deterministic models cannot match. In a world where the next supply chain disruption is always around the corner, the question isn't whether to adopt AI-driven optimization—it's how quickly you can get there.

Explore the full interactive dashboard and white paper

pharma-supply-chain-navigator

About the Author

Ilan Gleiser is a quantitative modeler and simulation architect specializing in agent-based models for supply chain resilience, financial risk, and geopolitical scenario

analysis. His work combines Monte Carlo methods with AI-driven optimization to deliver actionable intelligence for decision-makers in pharmaceuticals and financial services.