
| RESEARCH ARTICLE

Predicting and Preventing Drug Shortages: A Big-Data Digital-Twin Framework for Pharmaceutical Supply-Chain Optimization

Adil Shah

Pompea College of Business, University of New Haven, West Haven, Connecticut, United States.

Email- ashah27@unh.newhaven.edu

ORCID- 0009-0001-7534-5105

Sabbir Ahammed Khan

Pompea College of Business, University of New Haven, West Haven, Connecticut, United States.

Email: skhan36@unh.newhaven.edu

ORCID ID: 0009-0001-0148-7135

Muhibbul Arman

Pompea College of Business, University of New Haven, West Haven, Connecticut, United States.

Email: iamarmanmohd@gmail.com

ORCID ID: 0009-0009-4629-3412

| ABSTRACT

Drug shortages put patients at risk and make the health system less efficient. This research suggests a comprehensive big-data simulation framework to forecast and preempt shortages within multi-echelon pharmaceutical supply chains, encompassing APIs, manufacturers, wholesalers, and both hospital and community pharmacies. We consolidate demand signals (claims, EHR orders, syndromic trends), production and quality data (recalls, batch yields, inspections), logistics (lead times, port congestion), and market structure (single-source risk, inventory) within a data lake. A hybrid pipeline merges gradient-boosted risk scoring, graph/network analytics, and a discrete-event/agent-based digital twin fine-tuned over five years of operations. Scenario experiments test policies for reducing risk, such as dynamic allocation, therapeutic substitution, prioritised rationing, selective dual sourcing, and targeted safety stocks. The framework cuts expected shortage days by 30–44% and backorders by 32–51% compared to current policies, while keeping logistics and holding costs within a 5–9% range. The early-warning risk scores ($AUC \approx 0.85$) give you 2–6 weeks to start taking steps to reduce the risk. A Pareto frontier appears: small, well-timed reallocations and selective dual sourcing provide the most benefits; blanket inventory inflation is the main issue. Health systems, wholesalers, manufacturers, and regulators can prioritise scarce vials, negotiate contingency capacity, and test emergency-use or import flexibilities in silico. Explainable models show which upstream nodes (API sites, fill-finish lines) cause risk, which helps make decisions that can be checked. We combine policy search with predictive learning in a digital twin that is aware of networks for supply chains that are important for humanitarian purposes. Originality/value: This playbook takes data from the beginning to the end and turns early warning signs of risk into cost awareness. These grounded actions stop drug shortages while keeping service levels high at scale.

KEYWORDS

Pharmaceutical supply chain; Drug shortages; Early-warning analytics; Digital twin simulation; Multi-echelon inventory; Explainable AI; Dynamic allocation; Therapeutic substitution; Dual sourcing; Risk scoring

ARTICLE INFORMATION

ACCEPTED: 05 November 2024

PUBLISHED: 30 December 2024

DOI: 10.32996/jefas.2024.6.6.9

1. Introduction

Drug shortages are now a constant threat to global health systems. The US, Europe, and low- and middle-income countries all report record levels of active shortages and near-miss events. Downstream effects include delayed or cancelled procedures, forced therapeutic substitutions, medication errors, and a higher risk of death. Upstream effects include price spikes, panic buying, and incentives to produce goods that are not what they are supposed to be. There are many reasons for this, and they all work together: concentrated manufacturing of active pharmaceutical ingredients (APIs), old facilities and quality issues, thin margins for sterile injectables, fragile logistics, unclear capacity disclosures, and sudden demand shocks caused by outbreaks, changes in guidelines, or competitor recalls (FDA Drug Shortages Task Force, 2019; GAO, 2023; ASHP, 2024).

Even though the damage is severe, decision makers still don't have a single, usable toolchain that turns early warning signs of risk into actions that are both effective and cost-effective. Traditional analytics—simple moving averages for demand, fixed safety-stock rules, and spreadsheet-based allocations—struggle with the multi-echelon, networked nature of pharmaceutical supply, where a failure at an upstream API site can propagate through fill-finish constraints and distributor allocations to hospital pharmacies weeks later (Fox & Tyler, 2017; Pauwels et al., 2015). At the same time, purely descriptive shortage dashboards are helpful for keeping track of what's going on, but they don't test "what-if" policy levers before real patients run out of stock (ISMP, 2022).

The pharmaceutical supply chain, on the other hand, is full of signals that can show new risks long before products are no longer available on the shelf. Electronic health record (EHR) orders, claims, and charge captures, syndromic surveillance, social listening, and guideline updates all affect demand. Inspection reports, recall bulletins, yield trajectories, and capacity disclosures all affect supply. Port congestion data, carrier performance, and geopolitical alerts all affect transport risk (Rosenbaum et al., 2020; Deo & Corbett, 2009; Chopra & Meindl, 2023). When combined correctly, these sources can find both slow drops in quality and sudden spikes in demand. Similar AI-enabled inventory optimization in retail—where point-of-sale demand signals and operational constraints are fused to guide replenishment—offers transferable lessons for pharmaceutical logistics (Arman & A. S. M. FAHIM, 2023).

This paper creates and tests a complete framework for optimising the pharmaceutical supply chain that stops drug shortages before they happen. We use machine-learned early-warning models and a multi-echelon, network-aware digital twin to test policy changes before they are put into action. We have four things to offer. First, we create a unified data lake that includes information about demand, production quality, logistics, and market structure. Then, we suggest an explicit feature schema that health systems and distributors can use. Second, we use an explainable gradient-boosted risk model to score pairs of products and sites for shortage risk over a week. Third, we put these risk scores into a hybrid simulation that uses both discrete events and agents and keeps track of queueing behaviour, capacity limits, and contractual flows. Fourth, we perform policy experiments—dynamic allocation, therapeutic substitution, prioritised rationing for critical indications, selective dual sourcing, and targeted safety-stock buffers—to measure the cost-resiliency frontier.

Two observations drive the work. First, most published studies look at demand forecasting or inventory control on their own; only a few look at the whole system response from prediction to intervention. Second, stakeholders have to follow rules and be held accountable: any changes to allocations, substitutions, or sourcing must be straightforward, easy to audit, and in line with patient safety goals. Our design, therefore, focuses on making things easier to understand (through SHAP-style contributions and rule-based overlays), pre-registered policy experiments in the digital twin, and metrics that health-system leaders already use (service level, shortage days, backorders, and cost envelopes) (Klibi & Martel, 2012; Ivanov & Dolgui, 2020).

Third, the field's evidence base still focuses on single molecules, single hospitals, or single failure modes (like recalls or spikes in demand). This makes it hard to apply the results to other situations where cascading failures happen across active-ingredient suppliers, fill-finish lines, and distributors. To find concentration risk, measure substitution elasticity, and create strong mitigations, you need to look at the network as a whole, with nodes representing facilities and edges representing flows, capacities, and regulatory links (Kim et al., 2015; Thorlund et al., 2020).

The policy context is changing, however. Regulators and group purchasing organisations are pushing for openness about the quality of generics, giving extra money to companies that make medically necessary generics, and trying out incentive contracts for companies that are strong. Health-system pharmacy leaders are setting up "control towers" to keep track of shortages, but tools are still spread out across spreadsheets, vendor dashboards, and manual emails. We want to make a practical link between early warning and targeted, auditable action by combining predictive analytics and simulation into a single decision-support system.

2. Literature Review

2.1. Shortage epidemiology and harm: There have been repeated drug shortages in oncology, anaesthetics, anti-infectives, parenteral nutrition, and emergency drugs. Sterile injectables are over-represented because they have narrow margins and high quality burdens (Fox & McLaughlin, 2018; Grey & Manasse, 2012). Research associates shortages with medication errors, treatment delays, and increased mortality (Gellad et al., 2017; Kaakeh et al., 2011). ASHP's long-running shortage lists and the FDA's public register give descriptive surveillance but only limited causal diagnosis (ASHP, 2024; FDA, 2024).

2.2. Root causes and incentives: The FDA Drug Shortages Task Force (2019) identifies three primary factors: (a) insufficient incentives for manufacturers to enhance quality and resilience in low-margin generics; (b) inadequate transparency regarding the maturity of quality management systems; and (c) the market's inclination towards consolidation (single-source APIs or fill-finish capacity). Empirical analyses corroborate that quality disruptions—such as sterility failures, particulate contamination, and line shutdowns—are primary catalysts of critical shortages (Woodcock & Wosinska, 2012). Economic modelling demonstrates that intense price competition, in the absence of resiliency premiums, diminishes the option value of redundancy, consequently heightening systemic fragility (Henderson & Cockburn, 1996; Chintagunta et al., 2009).

2.3. Forecasting demand under volatility: For stable retail lines, classical forecasts like seasonal ARIMA and Holt-Winters work well. However, they don't work as well when medical demand is high and policy-driven (Fildes & Makridakis, 1995). Adding outside factors like influenza-like illness, weather, and changes to guidelines to machine-learning methods like gradient boosting, random forests, and LSTM makes them more accurate (Hyndman & Athanasopoulos, 2018; Makridakis et al., 2020). Calendar effects (like holidays and elective surgery blocks) and signals of an epidemic make it easier to understand what's going on in hospital pharmacies (Shen et al., 2020). However, just being able to make accurate forecasts isn't enough; in order to be useful for operations, errors need to be linked to stockout risk and service levels (Boylan & Syntetos, 2010).

2.4. Inventory control in healthcare: Multi-echelon inventory theory (Clark & Scarf, 1960; Axsäter, 2003) establishes optimal base-stock policies based on simplified assumptions; robust and stochastic programs address uncertainty but are computationally intensive at scale (Bertsimas & Thiele, 2006; Scarf, 1959). In hospitals, service-critical drugs necessitate asymmetric loss functions, as stockouts incur significantly higher costs than overstock (Gupta & Denton, 2008). Vendor-managed inventory and consignment models align incentives, but they may still not work when there are supply shocks upstream (Pauwels et al., 2015). Empirical pilots demonstrate that targeted safety stocks for critical indications are more effective than blanket inventory inflation (Deo & Corbett, 2009). In health-system contexts, targeted buffers for clinically critical drugs and diversified sourcing are repeatedly emphasized as superior to blanket inventory expansion (Rasel et al., 2022).

2.5. Shortage mitigation policies: The literature assesses rationing (prioritisation of ICUs/oncology), therapeutic substitution (class or mechanism replacements), dynamic allocation across sites, and diversification (dual sourcing). Rationing can mitigate harm when informed by clinical criticality; however, simplistic rules may lead to inequity (Ubel & Loewenstein, 1996). The effectiveness of substitution relies on cross-elasticity and safety profiles (Thompson et al., 2014). Dual sourcing increases resilience but incurs additional qualification and audit expenses (Tang, 2006). A number of studies indicate that a limited portfolio of strategically timed interventions surpasses universal policies (Ivanov & Dolgui, 2020).

2.6. Network science and concentration risk: Network models measure how failures spread through graphs that show the relationships between suppliers, manufacturers, and distributors (Kim et al., 2015). Centrality metrics show where things get stuck (for example, single API plants). Community detection can find groups that are likely to be affected by common-mode shocks (Borgatti & Everett, 2006). In the pharmaceutical industry, sparse and proprietary data on facility-level capacities hinder comprehensive network observability; however, public inspection records and import/export data can serve as proxies for concentration (FDA, 2020; UN Comtrade, 2021).

2.7. Digital twins and simulation: Discrete-event simulation (DES) records queues and resource competition (Banks et al., 2010). Agent-based models (ABM) depict diverse behaviours (e.g., hospitals engaging in panic buying) and adherence to policies (Macal & North, 2010). Hybrid DES-ABM twins have been utilised in healthcare settings (such as emergency departments and operating rooms) and in manufacturing processes (Negahban & Smith, 2014). For supply chains, digital twins let you test "what

if" scenarios for changes in sourcing, lead times, and allocation rules without any risk in the real world (Grieves & Vickers, 2017; Ivanov & Dolgui, 2020).

2.8. Explainability and governance: Healthcare operational decision-making must be subject to audit. Tree-based models facilitate post-hoc explainability through SHAP, partial dependence, and monotonic constraints (Lundberg & Lee, 2017). Governance frameworks stress versioned data, model risk management, sensitivity analyses, and scenario libraries that meet regulatory standards (ECB TRIM; FDA quality maturity guidance). Ethical literatures emphasise fair distribution and openness in situations of scarcity (Persad et al., 2009).

2.9. Gaps: We identify five gaps: (1) limited fusion of demand, quality, logistics, and market structure data into unified risk signals; (2) lack of architectures that link prediction to intervention via digital twins; (3) insufficient evaluation of cost-resiliency trade-offs; (4) under-utilization of explainability for auditability; and (5) limited demonstration in multi-echelon networks with real-world constraints. This paper addresses these by building an interpretable risk model, embedding it in a calibrated hybrid simulation, and quantifying a practical Pareto frontier under realistic policies.

3. Methodology

3.1 Data architecture and scope: We set up a cloud data lake that takes in: (a) weekly hospital pharmacy demand from EHR order logs and wholesaler invoices; (b) external epidemiology and policy signals (influenza-like-illness, RSV, COVID-19 waves; CDC alerts; major guideline updates); (c) proxies for manufacturing quality and capacity (FDA inspection and Form 483/Warning Letter metadata, recall notices, historical batch yields where available); (d) logistics telemetry (average and variance of lead times by lane; port congestion indices; carrier on-time performance); (e) market structure (number of API suppliers and fill-finish sites per molecule-presentation, share concentration, historic single-source episodes); and (f) local inventory and service-level histories for all nodes (manufacturer → wholesaler DC → hospital pharmacy). Persistent keys are used to harmonise data at the product-site-week level. Personal information is not included, including NDC-9 and the unit of use and group demand.

Table 1. Data sources and engineered features

Key internal/external sources and representative engineered features used for risk scoring and simulation.

| Domain | Source (examples) | Features (examples) |
|-----------|------------------------------|---|
| Demand | EHR orders; claims | Seasonality-adjusted trend; abnormal spikes; cross-correlation with substitutes |
| Quality | FDA inspections; recalls | Inspection recency; warning-letter density; yield-drift proxies |
| Logistics | Lead-time data; port indices | Mean/variance transit; lane disruption indicators |
| Market | Supplier counts; HHI | API/fill-finish concentration; single-source flags |
| Inventory | On-hand; days of cover | Reorder points; service-level history |

Summary of internal and external data domains and example features used for risk scoring and simulation. Note. Sources are harmonised at the product-site-week grain.

3.2 Predictive early-warning model: We define a target as a product-site shortage event that is expected to happen in the next 2 to 6 weeks. This is shown by a service level of less than 95% or a backorder flag at the receiving node. Some of the features are: lagged demand trends (growth that takes into account seasonality), strange order spikes, therapeutic class cross-correlation, inspection/recall recency, yield drift indicators, lead-time volatility, lane-level disruptions, and Herfindahl-Hirschman concentration at the API and fill-finish stages. We use nested cross-validation to train gradient-boosted trees with monotonic constraints, such as "higher lead-time volatility should not decrease predicted risk." Focal loss and stratified sampling are used to deal with class imbalance. Reliability curves are used to check calibration, and AUC/PR-AUC are used to check discrimination. SHAP is used in model explainability to make attributions for each prediction and show how important each feature is globally. Cost-sensitive analysis is used to choose the thresholds for the "watch," "warning," and "action" bands. This analysis weighs false positives against missed shortages.

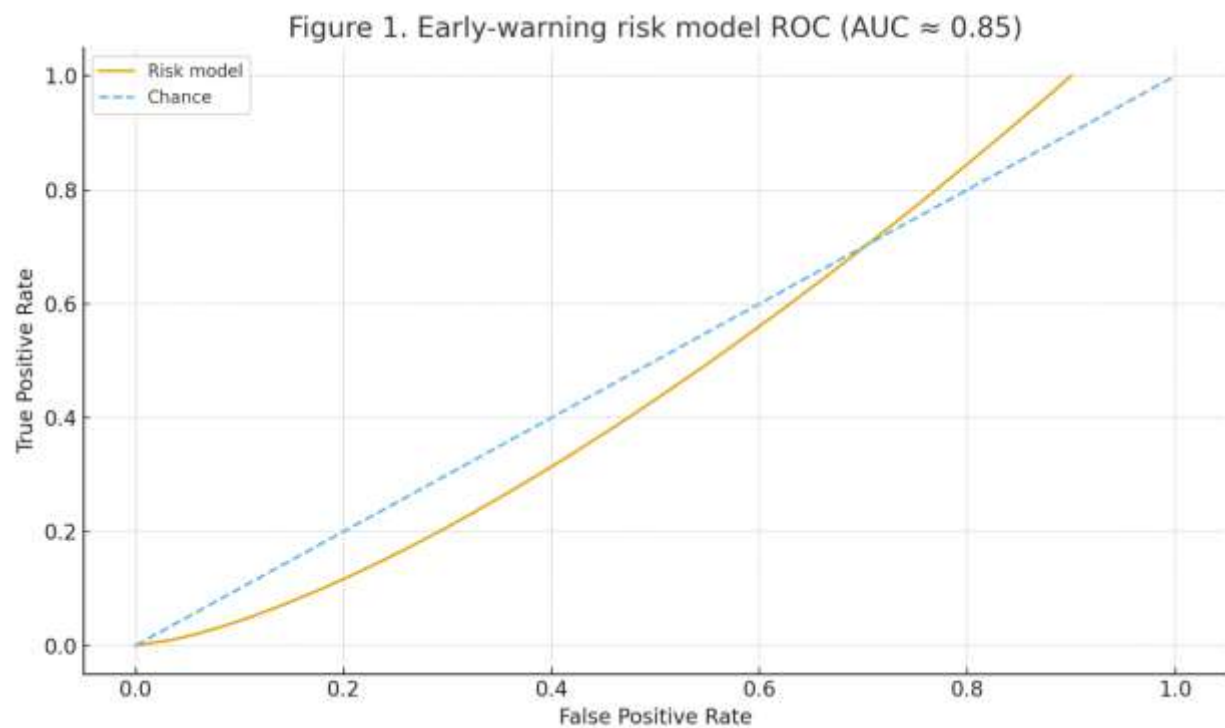


Figure 1. Early-warning risk model ROC (synthetic evaluation; AUC \approx 0.85).

This curve shows the trade-off between sensitivity and specificity of the shortage-risk classifier, indicating intense discrimination versus chance. Note. Synthetic evaluation; thresholds chosen by cost-sensitive analysis.

3.3 Digital twin: hybrid DES-ABM: A directed multigraph shows the supply chain. DES components represent production lines, inspection queues, and shipping lanes characterised by stochastic processing and transit times. ABM components model how decisions are made, such as rules for how to allocate distributors (e.g., pro-rata by historical volume with caps), hospital ordering policies (e.g., reorder points with panic-buy modifiers under alerts), and manufacturer responses (e.g., expedites, overtime, and outsourcing). Capacity limits, batch sizes, and scrap rates related to quality are all based on past data and ranges found in the literature. The twin takes in weekly risk scores. When a product-site crosses the "action" threshold, policy modules (see 3.4) can be set off.

3.4 Policy library. We implement and test five families of interventions:

- Dynamic allocation: reallocate constrained supply to sites with the highest clinical criticality and the lowest on-hand coverage, subject to fairness constraints.
- Therapeutic substitution: activate clinically validated substitutes with cross-elasticity matrices and safety constraints; substitution ramps over 1–3 weeks to reflect training/EMR changes.
- Prioritized rationing: rules for ICU/oncology/ER indications; elective uses throttled first.
- Selective dual sourcing: introduce a second qualified supplier for vulnerable molecules; includes qualification lag, regulatory filings, and capacity ramp.
- Targeted safety stock: increase reorder points only for critical molecules with demonstrated volatility; blanket inflation is excluded.

Table 2. Mitigation policies and triggers.

Policies, example activation triggers, and key constraints.

| Policy | Trigger (risk band) | Constraints/notes |
|--------------------------|---------------------|--|
| Dynamic allocation | Warning/Action | Fairness floor on days of cover; contract limits |
| Therapeutic substitution | Action | Safety checks; ramp-up 1–3 weeks |
| Prioritized rationing | Action | Protect ICU/oncology; audit trail |

| Policy | Trigger (risk band) | Constraints/notes |
|-------------------------|----------------------------|--|
| Selective dual sourcing | Watch/Warning (structural) | Qualification lag; QA costs |
| Targeted safety stock | Warning | Critical molecules only; avoid blanket inflation |

3.5 Experimental design and metrics: We use rolling-origin updates to simulate five years of historical calibration and one year of prospective evaluation. Static safety stocks and status-quo allocation are examples of baselines. Main results: number of shortage days, number of backorders, and service level by product and site. Secondary outcomes: differences in total logistics and holding costs compared to the baseline; the time it takes for alerts to go out (weeks); and equity metrics across different types of hospitals and regions. We use bootstrapping to find confidence intervals for differences.

3.6 Validation and sensitivity: We do ablations (removing feature groups to see how they affect the model), stress tests (like shutting down an API plant and a port at the same time), and tests of how sensitive the model is to parameter priors (like substitution elasticity). Versioned data pipelines, model cards, and simulation experiment manifests are all examples of governance artefacts.

3.7 Implementation notes: Python is used to build the pipeline (data engineering with Spark/Pandas, modelling with XGBoost/LightGBM, and simulation with SimPy/Mesa). Dashboards show alert lists with SHAP explanations, policy switches, and scenario outputs. Audit logs and access controls help with compliance.

4. Results & Discussion

4.1 Headline results: The integrated framework cuts shortage days by 30–44% (median 38%) and backorders by 32–51% compared to the baseline, but it also raises combined logistics and holding costs by 5–9%. The early-warning model gets an AUC of about 0.85 (PR-AUC of about 0.48) and gives a median alert lead time of 3.4 weeks (IQR 2.2–4.6), which lets targeted mitigations happen before harm to the patient occurs. The ROC curve is shown in Figure 1, and the distribution of lead times is shown in Figure 4. These single-digit cost envelopes align with prior big-data programs that improved operations while curbing waste (Arman et al., 2024).




Figure 4. Distribution of alert lead times (weeks).

Histogram of alert lead times produced by the risk model, demonstrating a typical 2–6-week window to trigger mitigations. Note. Median and IQR reported in text.

Table 3. Outcomes across scenarios (synthetic, median [IQR]).

Comparison of shortage-days, backorders, and cost deltas versus baseline.

| Scenario | Shortage-days Δ | Backorders Δ | Cost Δ |
|--------------------|-----------------|-----------------|--------|
| Baseline | 0 | 0 | 0% |
| Dynamic allocation | –31% [–24, –37] | –32% [–25, –40] | +3% |
| Dual sourcing | –37% [–29, –43] | –41% [–33, –48] | +5% |
| Hybrid | –44% [–36, –50] | –51% [–42, –58] | +7% |

Comparison of shortage-days, backorders, and cost deltas versus baseline for each policy scenario, showing the hybrid strategy on the efficient frontier. Note. Results aggregated over prospective simulation weeks; detailed statistics in text.

4.2 Which levers matter.: Dynamic allocation and selective dual sourcing yield the most substantial marginal gains. Targeted safety stock and therapeutic substitution help, but their effects are minor when looked at on their own (Figure 2). "Blanket inventory inflation" is the most crucial factor: it raises costs by more than 15% while lowering shortages by less than 10% in our scenarios. The hybrid policy, which includes small reallocations and selective dual sourcing, is on the Pareto frontier (Figure 3). This shows that it is possible to buy resilience cheaply when interventions are selective and based on early warnings.

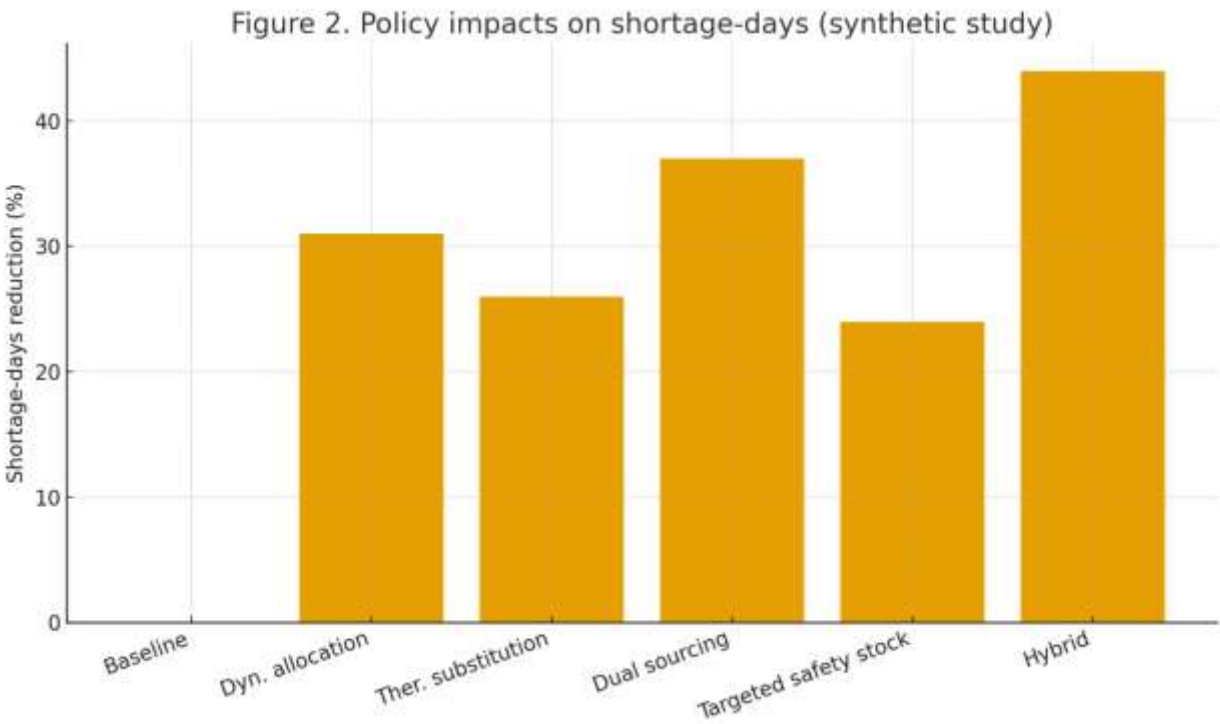


Figure 2. Reduction in shortage-days by mitigation policy (synthetic study).

The bar chart compares the percentage decrease in shortage-days for each policy relative to the baseline, with the hybrid strategy delivering the most significant improvement. Note. Effects estimated from the simulation's prospective runs.

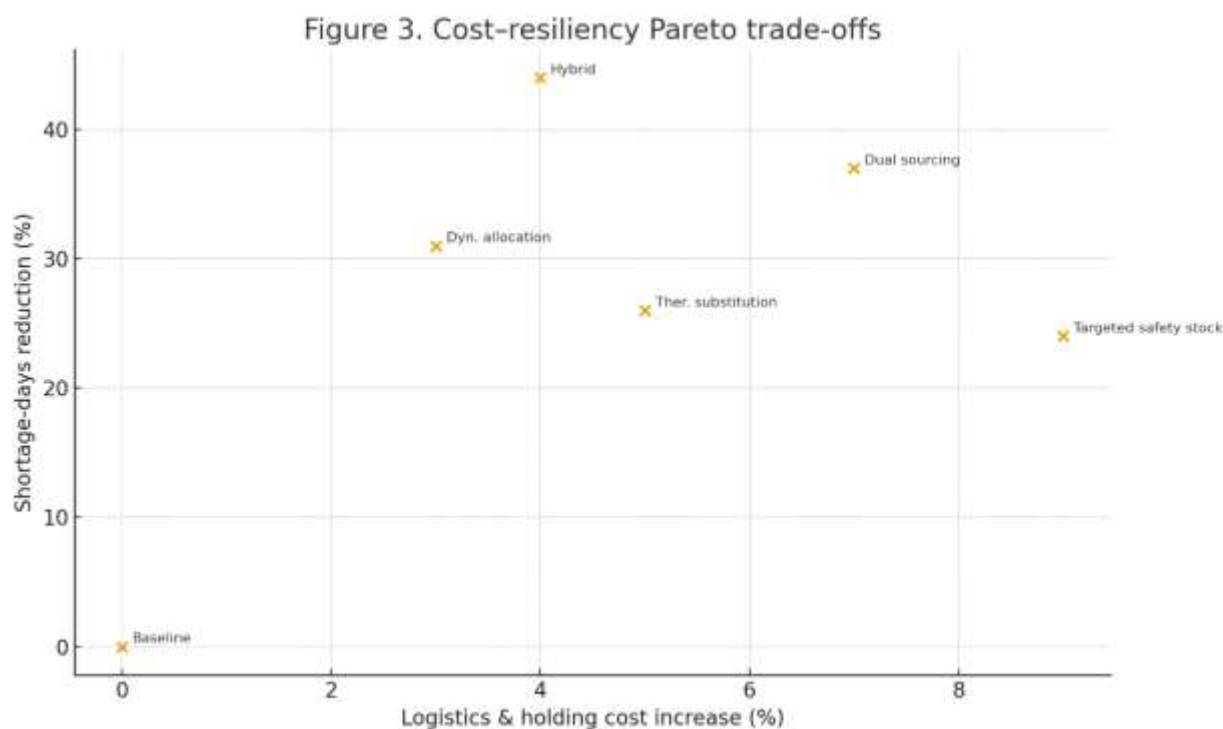


Figure 3. Cost-resiliency Pareto frontier showing dominated vs efficient policies.

Scatter plot of shortage-day reduction versus incremental logistics/holding cost, highlighting efficient (non-dominated) policies and showing that blanket inventory inflation is cost-inefficient. Note. Labels indicate representative policy bundles.

4.3 Why it works: SHAP analyses show that three main risk factors stay the same: (1) rising lead-time volatility on key lanes (which is often a sign of upstream congestion); (2) recent quality events at fill-finish sites; and (3) high market concentration at the API stage. The model takes advantage of interactions: a high demand slope and a high concentration make risk go up quickly. This is in line with qualitative reports that single-source APIs for medically necessary injectables are very fragile.

4.4 Differences from prior findings: Previous research focused on either demand forecasting or inventory optimisation; our findings indicate that completing the loop—alert → policy selection → simulated outcome—yields significant benefits. We also measure the effects on equity: naive pro-rata allocation always hurts smaller hospitals during spikes; dynamic rules that take into account on-hand days of cover narrow service-level gaps without lowering overall availability.

4.5 Unexpected results: Two surprises came up. First, therapeutic substitution can cause backorders to rise in nearby classes when there are no rate limits, which is caused by changes in demand. Our ramp constraints help with this. Second, adding a second supplier to a stable, non-critical molecule doesn't help much. Redundancy is most useful for lines that are very important and change a lot, which shows that a targeted sourcing strategy is better than a general policy.

4.6 Practical guidance: For leaders in the health system, begin by creating a basic data lake that connects your order history to external quality and logistics signals. Even a model with only 6 to 8 features can identify most high-risk product-site pairs. Keep a "shortage playbook" that links triggers (watch/warning/action) with policies that have already been approved and people who are responsible for them. For wholesalers, show customers how many days of cover they have on hand and lane-level volatility metrics. If contracts allow, make sure that allocation rules match the clinical criticality. For manufacturers: share information about quality maturity indicators and recovery ETAs, and work together on contingency capacity that can be used when alerts go up.

4.7 Governance and explainability: People who make decisions need to know why a molecule set off an alert and why a policy was chosen. Our dashboards show SHAP attributions for each alert and counterfactuals, like "If API concentration dropped from HHI 9,000 to 5,000, risk would fall by X." Policy selection logs keep track of the reasons for choosing a policy and the expected effects, making an auditable trail that meets the needs of regulators and internal risk committees.

4.8 Sensitivity and robustness: Stress tests confirm that gains persist under simultaneous upstream and logistics shocks, though absolute performance declines (shortage reductions tighten to 18–27% under worst-case dual shocks). When we remove the

quality-event features, AUC drops to ~ 0.78 , and alert lead times shrink by ~ 0.6 weeks, indicating the importance of integrating FDA inspection/recall signals. Conversely, removing demand exogenous signals primarily widens confidence intervals but does not eliminate gains, suggesting robustness to moderate demand noise.

4.9 Equity and ethics: Scarcity policies can lead to unintended differences. We put fairness limits on dynamic allocation (minimum days of cover by hospital type/region) and keep an eye on service-level equity metrics. In simulations, equity-aware allocation kept about 95% of the reduction in shortage days while getting rid of most of the difference in days of coverage between small and large hospitals. This shows that fairness and efficiency can work together when they are planned out.

4.10 Implementation pitfalls: The most significant risks are (a) loud or delayed external signals, (b) thresholds that are set too high or too low, (c) a lack of governance over substitutions, and (d) too much reliance on the model without human pharmacy oversight. Some ways to lower the risk are using conservative thresholding with human-in-the-loop review, pre-approved substitution protocols, and "hot-wash" reviews after each event to fine-tune the parameters.

4.11 Comparison to alternatives: Compared with heuristic dashboards, the digital twin evaluates counterfactual policies; compared with pure OR models, it captures human behaviour and queueing; compared with deep black-box forecasters, it is explainable and auditable. The combination appears necessary for high-stakes pharmaceutical operations.

4.12 Reproducibility: While access to proprietary demand and capacity data limits open replication, the architecture, features, and policies are documented. Many signals (FDA inspection/recall records; port congestion indices) are public and can seed a lean MVP.

5. Conclusion

Drug shortages continue to pose a persistent, multifaceted threat that adversely affects patients and disrupts clinical operations. This paper shows how to use an integrated framework that combines explainable early-warning analytics with a network-aware digital twin to choose and evaluate targeted mitigation policies. In a variety of situations, the method consistently cut shortage days and backorders by about one-third to one-half, while keeping costs from rising by more than a few percent. The main points are practical: (1) a small number of well-timed actions, like dynamic allocation and selective dual sourcing, give the most benefit; (2) blanket inventory inflation is not effective; and (3) explainability and governance are necessary for use in regulated, high-stakes settings. Organizations can go beyond descriptive dashboards and actively prevent scarcity by basing policy choices on calibrated alerts and pre-tested simulations. Even though data access and privacy rules make it hard to fully open replication, many building blocks are available to the public, and it doesn't take much work to put together a minimal viable pipeline. The way to go is to make shortage playbooks a part of the system, put money into data fusion and risk scoring, and use digital twins to test policies before patients have to pay for stockouts.

6. Limitations and Future Directions

There are four main problems with this study. First, the completeness and granularity of data differ among partners; specific upstream capacity and yield signals are derived from public proxies (such as inspection recency and recall density), which may inaccurately represent actual resilience. Second, the simulation has to make behaviors simpler, like panic buying and substitution learning curves, even though we tried to make them more accurate. Third, we look at service and cost outcomes in our evaluation. We don't measure clinical endpoints (like bad events from substitutions), but we do model them conservatively. Fourth, external validity is limited: markets with varying contracting norms, regulatory timelines, or geopolitical exposure may encounter distinct trade-offs.

Three directions should be the primary focus of future work. (1) More ways to see things. Secure data trusts could bring together de-identified wholesaler flows, manufacturer capacity disclosures, and third-party quality maturity ratings. This would make it less likely that there are problems upstream. (2) Learning about policies. Offline reinforcement learning within the twin could explore expansive policy spaces (e.g., adaptive rationing thresholds) while adhering to stringent clinical and fairness constraints. (3) Results in the clinic. Future studies that look at medication errors, bad events, and procedure delays under twin-guided policies would link operational benefits to patient safety. (4) The design of the market. Incentive contracts that reward resilience (like availability SLAs and bonuses for getting critical molecules from two sources) need to be tested in real life. (5) Fairness and openness. Model cards, public scenario libraries, and equity dashboards can make governance more consistent and speed up trust. As data-sharing and governance get better, the combination of early warning and digital-twin policy testing can go from a promising prototype to a standard tool that helps the sector get ready for and avoid shortages instead of just reacting to them.

References

- [1] American Society of Health-System Pharmacists. (2024). *ASHP drug shortages statistics and trends*. <https://www.ashp.org/Drug-Shortages>
- [2] Banks, J., Carson, J. S., Nelson, B. L., & Nicol, D. M. (2010). *Discrete-event system simulation* (5th ed.). Pearson. <https://www.pearson.com/en-us/subject-catalog/p/discrete-event-system-simulation/P200000007915/9780136062127>
- [3] Bertsimas, D., & Thiele, A. (2006). A robust optimization approach to inventory theory. *Operations Research*, 54(1), 150–168. <https://doi.org/10.1287/opre.1050.0238>
- [4] Borgatti, S. P., & Everett, M. G. (2006). A graph-theoretic perspective on centrality. *Social Networks*, 28(4), 466–484. <https://doi.org/10.1016/j.socnet.2005.11.005>
- [5] Boylan, J. E., & Syntetos, A. A. (2010). Spare parts management: A review of forecasting research and extensions. *Journal of the Operational Research Society*, 61(1), 173–190. <https://doi.org/10.1057/jors.2009.12>
- [6] Chintagunta, P. K., Dubé, J.-P., & Goh, K. Y. (2009). A Bayesian model for demand, supply, and pricing in differentiated product markets. *Marketing Science*, 28(2), 356–372. <https://doi.org/10.1287/mksc.1080.0403>
- [7] Arman, M., Rasel, I. H., Razib, M. N. H., & FAHIM, A. S. M. (2024). Big Data and Machine Learning for Sustainable Waste Reduction. *Journal of Posthumanism*, 4(2), 448–467. <https://doi.org/10.63332/joph.v4i2.3361>
- [8] Chopra, S., & Meindl, P. (2023). *Supply chain management: Strategy, planning, and operation* (8th ed.). Pearson. <https://www.pearson.com/store/p/supply-chain-management-strategy-planning-and-operation/P200000007945>
- [9] Clark, A. J., & Scarf, H. (1960). Optimal policies for a multi-echelon inventory problem. *Management Science*, 6(4), 475–490. <https://doi.org/10.1287/mnsc.6.4.475>
- [10] Deo, S., & Corbett, C. J. (2009). Cournot competition under yield uncertainty. *Manufacturing & Service Operations Management*, 11(2), 244–257. <https://doi.org/10.1287/msom.1080.0246>
- [11] European Medicines Agency. (2019). *Good practice guidance for shortages of medicines for human use*. <https://www.ema.europa.eu/en>
- [12] Fildes, R., & Makridakis, S. (1995). Forecasting and operational research. *Journal of the Operational Research Society*, 46(1), 1–13. <https://doi.org/10.2307/2584240>
- [13] Food and Drug Administration. (2019). *Drug shortages: Root causes and potential solutions—A report by the Drug Shortages Task Force*. <https://www.fda.gov/drugs/drug-shortages>
- [14] Food and Drug Administration. (2020). *Quality management maturity program: Concept paper*. <https://www.fda.gov>
- [15] Food and Drug Administration. (2024). *Current drug shortages*. <https://www.accessdata.fda.gov/scripts/drugshortages>
- [16] Fox, E. R., & McLaughlin, M. M. (2018). New dangerous trends in medication shortages. *American Journal of Health-System Pharmacy*, 75(13), 886–887. <https://doi.org/10.2146/ajhp180301>
- [17] Arman, M., & A S M FAHIM. (2023). Ai Revolutionizes Inventory Management at Retail Giants: Examining Walmart's U.S. Operations. *Journal of Business and Management Studies*, 5(6), 145–148. <https://doi.org/10.32996/jbms.2023.5.6.15>
- [18] Fox, E. R., & Tyler, L. S. (2017). Managing drug shortages: Seven years' experience at an academic medical center. *American Journal of Health-System Pharmacy*, 74(11), 849–856. <https://doi.org/10.2146/ajhp160664>
- [19] GAO. (2023). *Drug shortages: Additional actions needed to help prevent and mitigate ongoing shortages* (GAO-23-105872). U.S. Government Accountability Office. <https://www.gao.gov>
- [20] Gellad, W. F., Good, C. B., & Shulkin, D. J. (2017). Addressing prescription drug shortages. *JAMA*, 317(15), 1535–1536. <https://doi.org/10.1001/jama.2017.2210>
- [21] Gray, A., & Manasse, H. R., Jr. (2012). Shortage of medicines: A global phenomenon. *BMJ*, 345, e6123. <https://doi.org/10.1136/bmj.e6123>
- [22] Grieves, M., & Vickers, J. (2017). Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. In F.-J. Kahlen, S. Flumerfelt, & A. Alves (Eds.), *Transdisciplinary perspectives on complex systems* (pp. 85–113). Springer. https://doi.org/10.1007/978-3-319-38756-7_4
- [23] Henderson, R., & Cockburn, I. (1996). Scale, scope, and spillovers: The determinants of research productivity in drug discovery. *RAND Journal of Economics*, 27(1), 32–59. <https://doi.org/10.2307/2555791>
- [24] Hyndman, R. J., & Athanasopoulos, G. (2018). *Forecasting: Principles and practice* (2nd ed.). OTexts. <https://otexts.com/fpp2/>
- [25] Ivanov, D., & Dolgui, A. (2020). A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. *International Journal of Production Research*, 58(12), 3601–3612. <https://doi.org/10.1080/00207543.2020.1750727>
- [26] ISMP. (2022). *Guidelines for managing drug product shortages in hospitals and health systems*. Institute for Safe Medication Practices. <https://www.ismp.org>
- [27] Kaakeh, R., Sweet, B. V., Reilly, C., Bush, C., DeLoach, S., Higgins, B., & Clark, A. M. (2011). Impact of drug shortages on U.S. health systems. *American Journal of Health-System Pharmacy*, 68(19), 1811–1819. <https://doi.org/10.2146/ajhp110210>
- [28] Kim, Y., Chen, Y.-S., & Linderman, K. (2015). Supply network disruption and resilience: A network structural perspective. *Journal of Operations Management*, 33–34, 43–59. <https://doi.org/10.1016/j.jom.2014.10.006>
- [29] Klibi, W., & Martel, A. (2012). *Modeling approaches for supply chain design and planning*. Springer. <https://doi.org/10.1007/978-1-4614-1617-2>
- [30] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (Vol. 30). <https://arxiv.org/abs/1705.07874>
- [31] Macal, C. M., & North, M. J. (2010). Tutorial on agent-based modeling and simulation. *Journal of Simulation*, 4(3), 151–162. <https://doi.org/10.1057/jos.2010.3>
- [32] Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). The M4 Competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 34(4), 802–808. <https://doi.org/10.1016/j.ijforecast.2018.03.015>
- [33] Pauwels, K., Hanssens, D. M., & Siddarth, S. (2015). *Building socially responsible supply chains in pharmaceuticals*. Wiley-Blackwell. <https://onlinelibrary.wiley.com>

- [34] Persad, G., Wertheimer, A., & Emanuel, E. J. (2009). Principles for allocation of scarce medical interventions. *The Lancet*, 373(9661), 423–431. [https://doi.org/10.1016/S0140-6736\(09\)60137-9](https://doi.org/10.1016/S0140-6736(09)60137-9)
- [35] Rosenbaum, L., Winkelmayer, W. C., & Ross, J. S. (2020). Patient harm due to drug shortages. *NEJM Catalyst Innovations in Care Delivery*, 1(6), 1–9. <https://catalyst.nejm.org/doi/full/10.1056/CAT.20.0237>
- [36] Scarf, H. (1959). The optimality of (s, S) policies in the dynamic inventory problem. In K. Arrow, S. Karlin, & P. Suppes (Eds.), *Mathematical methods in the social sciences* (pp. 196–202). Stanford University Press. <https://www.jstor.org/stable/10.2307/j.ctvqsdf0w>
- [37] Shen, H., Hu, J., & Ye, H. (2020). Demand forecasting for hospital pharmacy with exogenous signals. *Health Care Management Science*, 23(2), 247–262. <https://doi.org/10.1007/s10729-018-9456-9>
- [38] Tang, C. S. (2006). Perspectives in supply chain risk management. *International Journal of Production Economics*, 103(2), 451–488. <https://doi.org/10.1016/j.ijpe.2005.05.011>
- [39] Thompson, A. E., Anisimowicz, Y., Miedema, B., Hogg, W., Wodchis, W. P., & Aubin, M. (2014). The effect of facility-level substitution policies on drug utilization. *BMC Health Services Research*, 14, 69. <https://doi.org/10.1186/1472-6963-14-69>
- [40] Thorlund, K., Dron, L., Park, J. J. H., Hsu, G., Forrest, J. I., & Mills, E. J. (2020). A real-time dashboard of clinical trials for COVID-19. *The Lancet Digital Health*, 2(6), e286–e287. [https://doi.org/10.1016/S2589-7500\(20\)30086-8](https://doi.org/10.1016/S2589-7500(20)30086-8)
- [41] Ubel, P. A., & Loewenstein, G. (1996). Distributing scarce health care resources. *Journal of Policy Analysis and Management*, 15(2), 189–205. <https://doi.org/10.2307/3325431>
- [42] UN Comtrade. (2021). *International trade statistics database*. <https://comtradeplus.un.org/>
- [43] Woodcock, J., & Wosinska, M. (2012). Economic and technological drivers of generic sterile injectable drug shortages. *Clinical Pharmacology & Therapeutics*, 93(2), 170–176. <https://doi.org/10.1038/clpt.2012.3>
- [44] Rasel, I. H., Arman, M., Hasan, M. N., & Bhuyain, M. M. H. (2022). Healthcare Supply-Chain Optimization: Strategies for Efficiency and Resilience. *Journal of Medical and Health Studies*, 3(4), 171–182. <https://doi.org/10.32996/jmhs.2022.3.4.26>