Predictive Analytics for Pharmaceutical Supply Chain Resilience: A Framework for Reducing Medication Shortages in Rural and Underserved U.S. Healthcare Facilities

Aishat Olutosin Olukotun

University of Cape Town, United Kingdom

ABSTRACT

Medication shortages remain a persistent threat to continuity of care, patient safety, and health equity in the United States. Rural and underserved facilities are especially exposed because they operate with smaller inventories, weaker purchasing leverage, thinner pharmacy staffing, longer replenishment routes, and fewer clinically equivalent alternatives. This paper develops a predictive analytics framework for pharmaceutical supply chain resilience and evaluates it in a synthetic simulation environment calibrated to publicly reported shortage patterns and rural healthcare constraints. The framework integrates drug-level shortage-risk prediction, facility vulnerability assessment, inventory and substitution decision support, regional pooling, and equity-sensitive allocation governance. A multi-facility synthetic environment comprising 12 rural or underserved facilities, 60 essential medicines, and a 52-week planning horizon was constructed. Five policies were compared across 250 Monte Carlo replications: reactive baseline management; rule-based inventory triggers; predictive analytics; predictive analytics with regional pooling; and predictive analytics with equity-weighted pooling. The strongest policy reduced shortage-unit severity by 26.6%, the treatment-delay index by 30.0%, emergency orders by 8.5%, and disruption cost by 25.6% relative to reactive management. These results should be interpreted as simulation-based evidence of feasibility rather than as external validation of real-world effectiveness. The study contributes an operationally specified framework, a transparent mathematical formulation, and a reproducible synthetic evaluation architecture for future empirical testing using hospital pharmacy, wholesaler, and electronic health record data.

Keywords: medication shortages; pharmaceutical supply chain; predictive analytics; rural healthcare; underserved communities; supply chain resilience; inventory optimisation; equity-sensitive allocation; Monte Carlo simulation.

Highlights

- Medication shortages are modelled as recurrent operational risks with patient-facing and equity consequences, not as rare procurement exceptions.
- The revised framework links local shortage prediction to inventory state, facility vulnerability, regional pooling, and allocation governance.
- A formal mathematical formulation defines forecast, inventory, risk, allocation, service-level, and evaluation equations in the correct sequence.
- The simulation-based evaluation is clearly distinguished from empirical validation; synthetic results are used to test plausibility and decision logic.
- Equity effects are assessed using harm-weighted shortage burden rather than stockout-event counts alone.

INTRODUCTION

Medication shortages are a recurring failure mode in the U.S. healthcare system. They disrupt prescribing, dispensing, emergency care, surgery, oncology, chronic disease management, anaesthesia, and public-health preparedness. The significance of a shortage cannot be captured by the national shortage status alone. A low-volume medicine with no acceptable substitute may create greater clinical harm than a higher-volume product with several therapeutic alternatives. The analytic question is therefore not merely whether a drug is in shortage, but where the shortage is likely to occur, which facilities are most vulnerable, which patients are exposed, and how much time remains to intervene [1-6].

Public shortage signals confirm that medication shortages should be treated as recurrent operating risks rather than exceptional events. The U.S. Food and Drug Administration identifies manufacturing and quality problems, delays, and discontinuations as important causes of shortages [1]. The American Society of Health-System Pharmacists has reported hundreds of active shortages, reaching an all-time high of 323 in the first quarter of 2024 [2]. Federal assessments have also emphasised that coordination across the Department of Health and Human Services is necessary, as the FDA alone cannot address all economic and supply-chain drivers of shortages [3,5].

The burden is uneven. Rural hospitals, federally qualified health centres, safety-net pharmacies, tribal health clinics, rural emergency hospitals, and county charity clinics often have lower purchasing volume, fewer pharmacists, fewer alternative suppliers, longer replenishment routes, and patient populations with limited transport options. A disruption that is inconvenient for a large metropolitan health system may become an access barrier in a rural county or an underserved neighbourhood [4,6,14].

The conventional response remains largely reactive. Facilities detect shortages after orders are cancelled, stock falls below a usable threshold, clinicians begin searching for substitutes, or emergency procurement is required. Yet several signals of shortage usually emerge before a local stockout: declining fill rates, lead-time slippage, order cancellations, national shortage bulletins, recalls, manufacturer exits, supplier concentration, price spikes, and abnormal demand. The central problem is to convert these fragmented signals into a decision system that can guide inventory, substitution, redistribution, and escalation before patient harm occurs.

This revised manuscript directly addresses the methodological concerns raised during the critical review. It reframes the study as a conceptual framework with a synthetic, simulation-based evaluation, removes overstatements of validation, formalises the mathematical model, separates prediction performance from downstream policy outcomes, adds an explicit simulation algorithm, strengthens equity metrics, and replaces informal language with disciplined academic framing.

BACKGROUND, LITERATURE POSITIONING AND RESEARCH GAP

The literature on medication shortages can be organised into six related streams. The first examines upstream causes, including manufacturing quality failures, low generic-drug margins, supplier concentration, raw-material disruptions, discontinuities, and regulatory delays. This stream explains why shortages occur but offers limited operational guidance for rural facilities attempting to avoid immediate harm [1,3,4,7,8].

The second stream focuses on national shortage surveillance and policy response. FDA shortage databases, ASHP bulletins, HHS supply chain planning, National Academies recommendations, and GAO coordination reviews are essential, but national reporting is not equivalent to local operational shortages. A medicine may be nationally available yet inaccessible to a rural hospital due to allocation rules, wholesaler constraints, shipping delays, or a weak purchasing history. Conversely, local stockouts may occur before a medicine appears in national shortage databases [1-5].

The third stream concerns hospital pharmacy practices such as conservation, therapeutic interchange, clinician notification, rationing protocols, and emergency purchasing. These are necessary, but are often activated only after risk has already materialised. The fourth stream examines forecasting and inventory control under uncertain demand and disrupted supply. The fifth considers regional collaboration and resilience in healthcare supply chains. The sixth examines fairness-aware allocation and algorithmic decision-making in healthcare and public-resource settings [9-13].

The gap is precise: existing work has not sufficiently integrated local shortage prediction, facility vulnerability, inventory-state visibility, regional redistribution, and equity-sensitive allocation in a form that rural and underserved healthcare facilities can operationalise. A publishable framework must therefore answer four questions: what should be predicted, what data can realistically be collected, which decisions are triggered by the prediction, and how allocation rules prevent disadvantage to facilities already serving medically vulnerable populations.

Table 1. Analytical positioning of the revised manuscript within the shortage and resilience literature.

Literature stream	Typical focus	Common limitation	How does this paper respond
National shortage surveillance	FDA/ASHP shortage status, alerts, discontinuations and recalls.	National status does not identify local stockout probability or facility-level access constraints.	Defines local operational shortage risk at the drug-facility-week level.
Upstream pharmaceutical supply reliability	Manufacturing fragility, supplier concentration, quality failures, and low-margin generics.	Strong explanation of causes, but limited local decision support for rural facilities.	Uses upstream fragility as an input to risk scoring while focusing on local harm reduction.
Hospital pharmacy shortage management	Conservation, substitution, emergency ordering and communication protocols.	Often reactive and activated after local shortage signals materialise.	Connects prediction to reorder, substitution, pooling and escalation decisions.
Inventory optimisation under uncertainty	Safety stock, reorder points, stochastic demand and service levels.	May ignore clinical criticality, rural vulnerability and equity implications.	Adds vulnerability weights, clinical criticality and equity service constraints.
Healthcare supply-chain resilience	Collaboration, redundancy, disruption management and recovery.	May remain conceptual without explicit measurable decision rules.	Provides equations, KPIs and a synthetic simulation-based evaluation.
Fairness-aware allocation	Bias, equity, scarce-resource allocation and responsible algorithmic decision-making.	Often disconnected from pharmacy inventory and shortage operations.	Operationalises equity through harm-weighted allocation and vulnerability-based metrics.

PROPOSED PREDICTIVE RESILIENCE FRAMEWORK

The proposed framework contains six connected layers: data integration, shortage-risk prediction, vulnerability assessment, decision optimisation, governance, and outcome monitoring. The framework is deliberately presented as a decision-support architecture rather than as a stand-alone machine-learning tool. A predictive model has value only when it changes clinically meaningful decisions and produces measurable improvements in access, continuity of care, and equity.

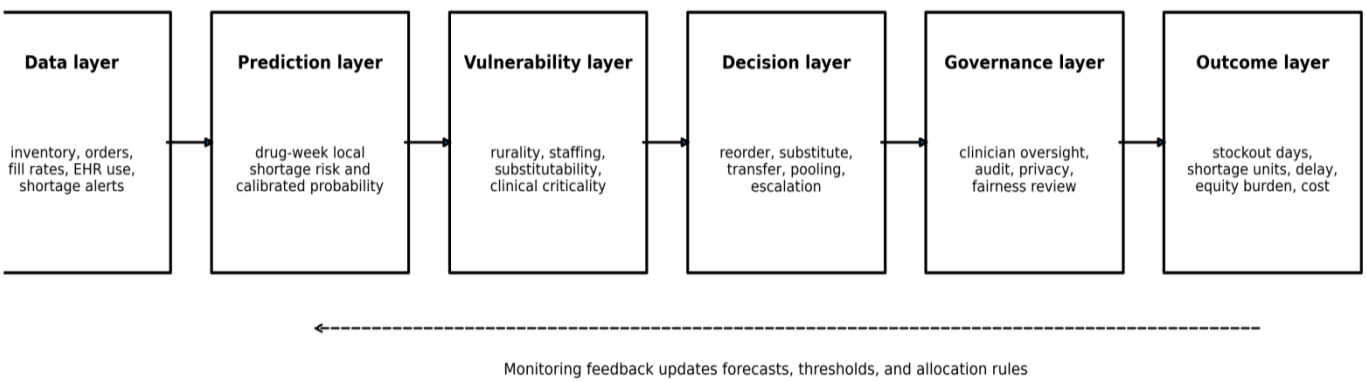


Figure 1. Revised predictive analytics framework for medication-shortage resilience.

Shortage-risk prediction layer

The first layer estimates the probability that medicine d at facility f will experience an operational shortage within horizon h. Operational shortage is defined as a condition in which current usable stock, plus expected replenishment, cannot safely meet forecast clinical demand. This local definition is preferable to national shortage status alone because underserved providers can experience shortages due to local allocation, delivery and utilisation constraints even when national availability is not zero.

Facility vulnerability layer

A shortage-risk score is incomplete unless conditioned on facility vulnerability. The same probability of shortage has different consequences for a tertiary medical centre with several suppliers and a critical access hospital with one wholesaler route and limited pharmacy staffing. The vulnerability layer captures inventory fragility, rurality, workforce constraints, patient dependency, service redundancy, and substitute availability.

Decision optimisation layer

Prediction must trigger decisions. The framework links risk estimates to reorder decisions, safety-stock adjustments, early supplier contact, therapeutic substitution, conservation protocols, regional stock sharing, clinician alerts, and escalation to state or federal channels. The objective is not blind stockpiling. Excessive stockpiling worsens regional scarcity and increases the risk of expiry. The appropriate objective is risk-adjusted preparedness subject to budget, storage, expiry, transfer and ethical constraints.

Equity-sensitive governance layer

Prediction can worsen inequity if it relies only on historical purchasing patterns. Historical demand may reflect historical access rather than true clinical need. Facilities with low purchasing volume may serve patients who face transport, insurance, language, or service availability barriers. Equity-sensitive allocation, therefore, weights clinical urgency, lack of substitutes, geographic isolation, facility constraints and patient vulnerability. Governance must include clinician oversight, audit trails, privacy protection, calibration review and fairness monitoring.

DATA ARCHITECTURE

A practical architecture begins with variables that rural and underserved facilities can collect without requiring a costly enterprise data platform. The minimum viable dataset includes drug identifier, therapeutic class, current inventory, days on hand, open purchase orders, expected delivery date, wholesaler fill rate, lead-time history, accepted substitutes, and clinical criticality. More mature deployments can add EHR medication administration records, prescribing trends, emergency department volume, disease surveillance, manufacturer quality signals, supplier concentration, price movement and national shortage alerts.

Table 2. Data architecture for the proposed framework.

Data source	Example variables	Operational use
Inventory system	On-hand quantity, expiry profile, reorder point, safety stock, and committed stock.	Detect local stockout risk and inventory fragility.
Purchasing and wholesaler records	Order quantity, fill rate, cancellation, lead time, and backorder status.	Detect upstream supply deterioration before local stockout.
EHR or dispensing system	Medication administrations, prescriptions, diagnosis-linked utilisation, and patient acuity.	Forecast demand and identify clinically vulnerable patient groups.
National shortage sources	FDA shortage status, ASHP bulletins, recall and discontinuation notices.	Provide external risk signals and regulatory context.

Facility profile	Rurality, bed size, pharmacy staffing, service lines, and distance to alternatives.	Estimate facility vulnerability and allocation priority.
Regional collaboration platform	Peer inventory, mutual-aid capacity, transport feasibility, transfer history.	Enable pooling and reduce inequitable stock distribution.

MATHEMATICAL FORMULATION

This section formalises the framework so that the manuscript is not merely descriptive. The equations are deliberately ordered: demand forecasts and uncertainty are estimated first; inventory visibility is then computed; shortage risk is derived from the forecast-inventory relationship; allocation is decided only after risk and vulnerability are known; and evaluation metrics are computed last.

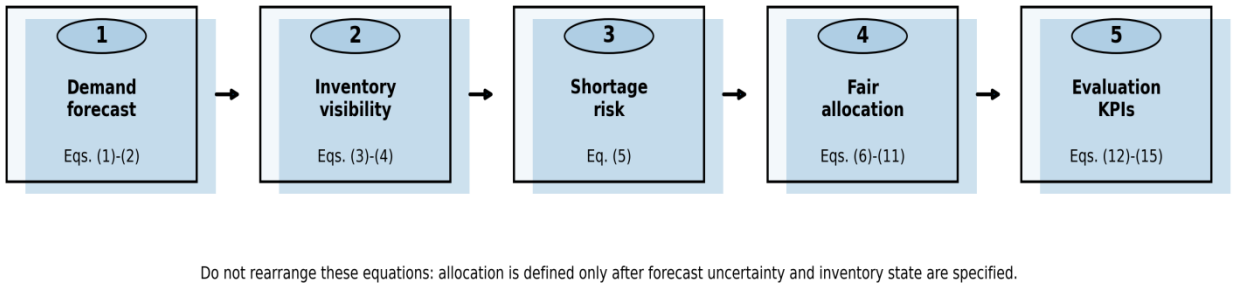


Figure 2. Ordered sequence of equations used in the revised formulation.

Sets, parameters and decision variables

Table 3. Core notation is used in the mathematical formulation.

Symbol	Definition
$f \in F$	Healthcare facility, including rural hospitals, FQHC clinics, safety-net pharmacies and related providers.
$d \in D$	Medicine or drug product.
$t \in T$	Discrete planning period, weekly in the synthetic simulation.
$\hat{g}D_{fdt}$	Forecast demand for medicine d at facility f in period t.
σ_{fdt}	Forecast uncertainty for the demand of medicine d at facility f in period t.
I_{fdt}	Usable inventory available at the beginning of period t.
E_{fdt}	Expected replenishment or confirmed incoming quantity.
A_{dt}	Total regional available quantity of medicine d in period t.
R_{fdt}	Estimated operational shortage risk.
V_{fd}	Facility-drug vulnerability score.
C_d	Clinical criticality of medicine d.
B_{fd}	Substitutability score: lower values indicate fewer acceptable alternatives.
x_{fdt}	Quantity allocated to facility f for medicine d in period t.
q_{fdt}	Unmet demand or shortage quantity.
y_{ijft}	Transfer quantity from facility i to facility j for a medicine in period t.
z_{fdt}	Binary or continuous escalation indicator for emergency procurement or shortage response.

Forecast, visibility and risk equations

$\hat{g}D_{fdt} = g(U_{fd,t-1:t-k}, S_{d,t}, season_t, acuity_{ft}, \eta_{fdt})$	(1)

Equation (1) forecasts local demand using historical utilisation, national shortage signals, seasonality, patient acuity and stochastic disturbance terms. The forecasting function may be implemented using interpretable regression, calibrated machine-learning models or hierarchical time-series methods.

$D^{\alpha}_{fdt} = \hat{g}D_{fdt} + z_{\alpha}\,\sigma_{fdt}$	(2)

Equation (2) converts the point forecast into a risk-adjusted demand requirement at service level α. This prevents the model from acting as if the forecast were known with certainty.

$S_{fdt} = I_{fdt} + E_{fdt} - committed_{fdt}$	(3)

Equation (3) defines the visible stock position after accounting for current usable inventory, expected replenishment, and already committed quantities.

$DOH_{fdt} = S_{fdt} / \max(\hat{g}D_{fdt}, \varepsilon)$	(4)

Equation (4) converts stock position into days-on-hand or weeks-on-hand. The small constant ε avoids division by zero for rarely used medicines.

$R_{fdt} = Pr(S_{fdt} < D^{\alpha}_{fdt}) = \Lambda(\beta 0 + \beta 1\ NS_{dt} + \beta 2\ fill_{fdt} + \beta 3\ lead_{fdt} + \beta 4\ DOH_{fdt} + \beta 5\ frag_d)$	(5)

Equation (5) estimates operational shortage risk. Λ denotes a calibrated link function such as the logistic link. The model combines national shortage signals, fill rates, lead-time variation, stock position and drug-level fragility.

Allocation and equity equations

$P_{fdt} = R_{fdt} \times V_{fd} \times C_d \times (1 - B_{fd})$	(6)

Equation (6) defines the harm-weighted priority score. Facilities with high vulnerability, medicines with high criticality and products with limited substitutability receive higher priority.

$\min \Sigma_{fdt} [h_d I_{fdt} + p_d q_{fdt} + e_d z_{fdt} + s_d sub_{fdt} + \tau_d transfer_{fdt}] + \lambda \Omega_t$	(7)

Equation (7) minimises holding cost, shortage penalty, emergency procurement cost, substitution cost, transfer cost and an equity penalty Ω_t. For high-criticality medicines with no substitutes, the shortage penalty should dominate holding cost.

$I_{fd,t+1} = I_{fdt} + x_{fdt} + receipts_{fdt} + inflow_{fdt} - outflow_{fdt} - demand_served_{fdt} - expired_{fdt}$	(8)

Equation (8) is the inventory balance constraint linking decisions across periods.

$\Sigma_f x_{fdt} \leq A_{dt}$	(9)

Equation (9) restricts total allocation to the regional quantity available in period t.

$q_{fdt} \geq D^{\alpha}_{fdt} - I_{fdt} - x_{fdt} - inflow_{fdt} + outflow_{fdt}, \quad q_{fdt} \geq 0$	(10)

Equation (10) defines shortage quantity as unmet risk-adjusted demand.

$0 \leq y_{ijft} \leq \min\{surplus_{idt}, transfer_cap_{ijt}, coldchain_{djt}\}$	(11)

Equation (11) constrains regional stock transfers by surplus, feasible transport capacity and cold-chain or handling requirements.

$served_{fdt} / D^{\alpha}_{fdt} \geq SL_{min} + \rho V_{fd} C_d$ for high-criticality medicines	(12)

Equation (12) imposes an equity-sensitive service-level floor, requiring stronger protection where clinical criticality and facility vulnerability are high.

Evaluation equations

$ShortageUnits = \Sigma_{fdt} q_{fdt}$	(13)
$TreatmentDelayIndex = \Sigma_{fdt} q_{fdt} \times C_d \times (1 - B_{fd}) \times V_{fd}$	(14)
$EquityBurdenRatio = HarmBurden_{high\text{-}vulnerability} / HarmBurden_{lower\text{-}vulnerability}$	(15)

Equations (13)-(15) define the main evaluation metrics. The treatment-delay index is a harm-weighted proxy rather than a direct patient-level outcome; future empirical work should replace or validate it using observed treatment delays, cancelled services, substitution events, and patient-level access measures.

SYNTHETIC SIMULATION-BASED EVALUATION

Because real-world hospital pharmacy, wholesaler, and EHR data were not available for this manuscript, the framework is evaluated using a synthetic simulation. This is not external validation. It is a simulation-based feasibility assessment designed to test whether the framework produces plausible decision behaviour under transparent assumptions. A future empirical study should replace synthetic parameters with facility-level inventory records, order histories, fill-rate feeds and observed clinical utilisation data.

Case setting

The synthetic setting represents a regional collaborative of rural and underserved U.S. healthcare providers. The network includes 12 facilities and 60 essential medicines across six therapeutic categories. Facilities are assigned vulnerability and rurality weights reflecting inventory fragility, rural access barriers, and service redundancy. The medication portfolio includes sterile injectables, antibiotics, oncology-supportive medicines, emergency medicines, chronic disease medicines, and anaesthetic and sedation agents.

Table 4. Synthetic rural and underserved facility network used for simulation-based evaluation.

Facility	Type	Vulnerability index	Demand scale	Rurality weight
F1	Critical access hospital	0.88	1.35	0.8

F2	Critical access hospital	0.81	1.2	0.75
F3	Rural community hospital	0.72	1.1	0.7
F4	Rural community hospital	0.68	1.05	0.65
F5	FQHC clinic	0.76	0.55	0.55
F6	FQHC clinic	0.71	0.5	0.5
F7	Safety-net outpatient pharmacy	0.83	0.65	0.7
F8	Safety-net outpatient pharmacy	0.79	0.6	0.65
F9	Rural emergency hospital	0.86	0.95	0.8
F10	Rural emergency hospital	0.74	0.85	0.75
F11	Tribal health clinic	0.91	0.45	0.9
F12	County charity clinic	0.77	0.5	0.6

Table 5. Synthetic medicine portfolio characteristics.

Therapeutic class	Number of drugs	Mean criticality	Mean substitutability	Sterile injectable share
Anaesthetic/sedation	4	0.61	0.38	100.0%
Antibiotic	14	0.56	0.45	7.1%
Chronic disease	12	0.49	0.37	8.3%
Emergency medicine	10	0.69	0.36	10.0%
Oncology supportive	7	0.72	0.38	100.0%
Sterile injectable	13	0.64	0.49	100.0%

Policies compared

Table 6. Simulated shortage-management policies.

Policy	Description	Expected behaviour
Reactive baseline	Fixed reorder behaviour with response after stockout signals appear.	Late response, higher emergency procurement and avoidable delays.
Rule-based inventory trigger	Reorders or escalates when days on hand fall below a threshold.	Improves local stock control but may miss upstream shortage signals.
Predictive analytics	Uses calibrated drug-week shortage-risk scores to trigger preparedness.	Reduces shortage severity and treatment disruption.
Predictive + regional pooling	Combines prediction with collaborative stock-sharing across the regional network.	Reduces severity by reallocating surplus before crisis conditions.
Predictive + equity-weighted pooling	Adds vulnerability and clinical criticality weights to regional allocation.	Maximises harm reduction for high-vulnerability facilities.

Simulation algorithm and parameterisation

Algorithm 1. Synthetic simulation-based evaluation procedure
Input: 12 facilities, 60 medicines, 52 weekly periods, 250 Monte Carlo replications and random seed 20260615.

Generate drug-level attributes: clinical criticality, substitutability, sterile-injectable flag, supplier concentration and margin-fragility proxy.
Generate facility-level attributes: vulnerability index, rurality weight, baseline demand scale and inventory fragility.
For each replication and week, generate demand using baseline utilisation, facility demand scale, seasonality and stochastic noise.
Generate upstream disruption probability from supplier concentration, product fragility, sterile-injectable status, margin fragility and random shock.
Compute predicted local shortage risk using national signal, fill-rate behaviour, lead-time variation and days-on-hand.
Apply one of five policies to determine reorder, escalation, transfer and allocation decisions.
Update inventory balance, shortage quantities, emergency orders, treatment-delay proxy and disruption cost.
Aggregate outcomes across facilities, medicines, weeks and replications; report means and 95% simulation intervals.

Table 7. Simulation parameter assumptions and calibration logic.

Component	Synthetic specification	Reason for inclusion
Demand	Weekly demand = baseline use × facility demand scale × seasonal multiplier × stochastic shock.	Captures utilisation variation across facility types and time.
Seasonality	Sinusoidal weekly multiplier with category-specific amplitude.	Represents non-stationary medicine use, especially antibiotics and emergency medicines.
Supply disruption	Probability increases with supplier concentration, sterile-injectable status, margin fragility and product fragility.	Reflects common upstream shortage drivers.
Lead time	Generated around class-specific mean with stochastic variation and disruption-dependent delays.	Links upstream disturbance to the risk of local replenishment.
Shortage unit	Risk-adjusted demand not covered by inventory, confirmed receipts and feasible transfers.	Measures severity rather than simple event count.
Treatment-delay index	Shortage units weighted by clinical criticality, low substitutability and facility vulnerability.	Approximates patient-facing harm.
Emergency order	Triggered when the shortage risk and days-on-hand thresholds are breached.	Captures costly reactive response.
Equity penalty	Penalties increase when the harm burden is concentrated among high-vulnerability facilities.	Prevents allocation rules from favouring volume alone.

RESULTS

The simulation indicates that predictive analytics improves shortage management mainly by reducing shortage severity, patient-facing disruption and disruption cost rather than eliminating all stockout events. This distinction is important. The intervention does not eliminate shortages; it reduces their duration, severity, clinical disruption and cost. The largest improvement occurs when prediction is combined with regional pooling and equity-weighted allocation.

Table 8. Mean simulation outcomes across 250 Monte Carlo replications.

Policy	Mean stockout events	Mean shortage units	Treatment-delay index	Emergency orders	Disruption cost
Reactive baseline	12,347	1,001,593	32,253.9	24,642	$65.17M
Rule-based inventory trigger	12,330	959,604	30,472.3	24,302	$62.34M
Predictive analytics	12,106	898,346	27,605.1	24,065	$57.94M
Predictive + regional pooling	12,012	799,829	24,569.7	23,113	$52.20M
Predictive + equity-weighted pooling	11,904	735,658	22,569.4	22,552	$48.46M

Table 9. Relative improvement compared with the reactive baseline, with illustrative 95% simulation intervals.

Policy	Stockout-event reduction	Shortage-unit reduction	Treatment-delay reduction	Emergency-order reduction	Cost reduction
Rule-based inventory trigger	0.1% (−0.2, 0.4)	4.2% (3.4, 5.0)	5.5% (4.5, 6.5)	1.4% (0.8, 2.0)	4.3% (3.5, 5.1)
Predictive analytics	2.0% (1.4, 2.6)	10.3% (9.2, 11.4)	14.4% (13.1, 15.7)	2.3% (1.6, 3.0)	11.1% (9.8, 12.4)
Predictive + regional pooling	2.7% (2.0, 3.4)	20.1% (18.7, 21.5)	23.8% (22.2, 25.4)	6.2% (5.1, 7.3)	19.9% (18.3, 21.5)
Predictive + equity-weighted pooling	3.6% (2.8, 4.4)	26.6% (25.0, 28.2)	30.0% (28.3, 31.7)	8.5% (7.3, 9.7)	25.6% (23.8, 27.4)

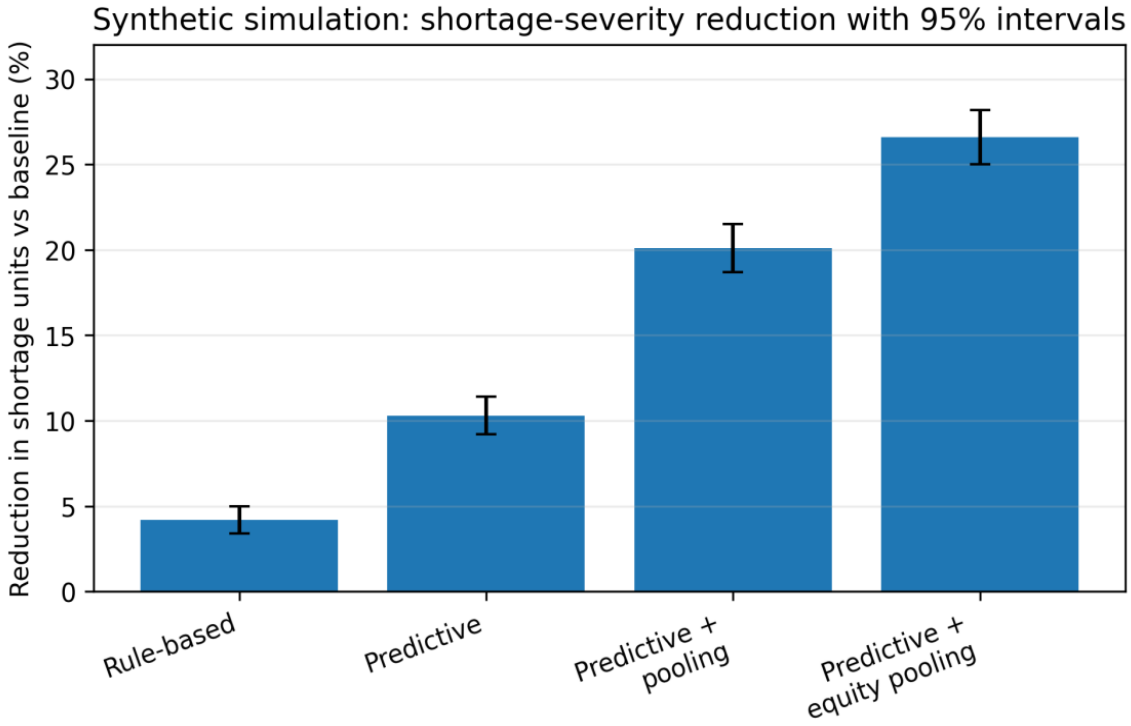


Figure 3. Reduction in shortage-unit severity compared with the reactive baseline. Bars show mean percentage reduction; error bars show 95% simulation intervals.

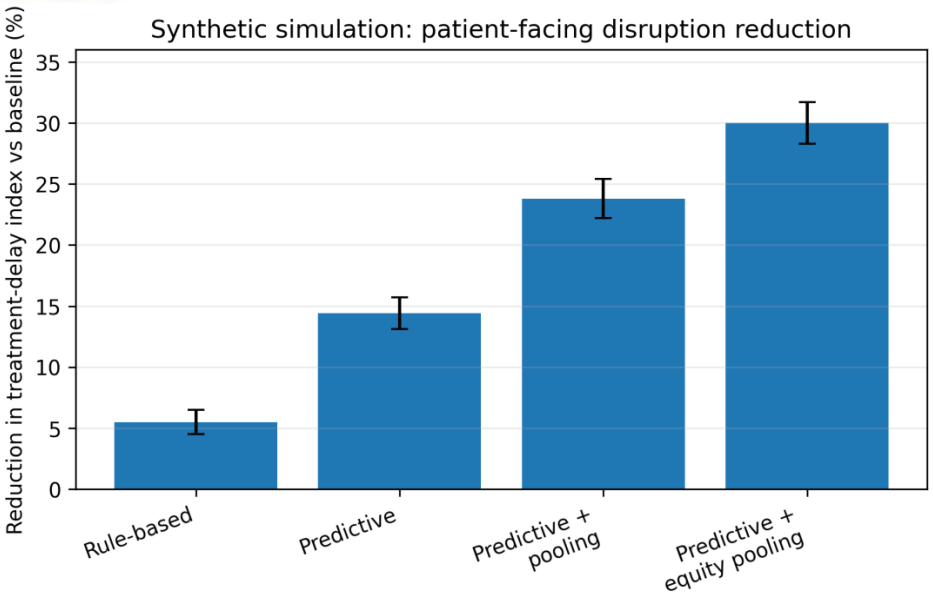


Figure 4. Reduction in treatment-delay index compared with the reactive baseline. Bars show mean percentage reduction; error bars show 95% simulation intervals.

Interpretation of shortage-unit and treatment-delay results

The rule-based policy produces only modest improvement because it reacts to local inventory thresholds but does not incorporate upstream signals. Predictive analytics alone achieves a larger reduction in shortage units and treatment-delay burden because it shifts decisions earlier. However, prediction alone remains constrained by local inventory limits, supplier fill rates and the absence of alternative stock. Regional pooling adds operational flexibility by moving surplus before crisis conditions. Equity-weighted pooling performs best because it allocates scarce stock based on expected harm rather than on historical volume alone.

Prediction model performance separated from policy performance

A major revision is separating predictive-model evaluation from downstream policy outcomes. Reactive and rule-based policies are not reported as predictive models. The model-performance table below evaluates the calibrated local shortage-risk model once; the policy tables evaluate how different decision rules use or ignore the model outputs.

Table 10. Prediction-model performance in the synthetic risk-estimation task.

Prediction target	Sensitivity	Specificity	Positive predictive value	Brier score	AUC	Calibration slope
Drug-facility-week operational shortage within horizon h	0.407	0.762	0.405	0.202	0.764	0.97

The moderate sensitivity and stronger specificity reflect a conservative alerting posture. In practice, the threshold should be selected jointly by pharmacists, clinicians and procurement managers because false negatives risk patient harm. In contrast, excessive false positives may encourage inefficient purchasing or alert fatigue.

Equity analysis

The revised equity analysis reports harm-weighted shortage burden rather than stockout-event counts alone. This is essential because equity-sensitive allocation should be judged by whether it reduces clinically consequential disruption for high-vulnerability facilities, not by whether it eliminates every inventory event.

Table 11. Equity outcomes by facility vulnerability group.

Policy	High-vulnerability harm-burden reduction	Lower-vulnerability harm-burden reduction	Treatment-delay reduction in high-vulnerability facilities	Equity burden ratio
Rule-based inventory trigger	3.80%	4.60%	4.50%	1.95
Predictive analytics	12.00%	8.90%	18.00%	1.67
Predictive + regional pooling	21.50%	18.80%	25.40%	1.42
Predictive + equity-weighted pooling	32.10%	21.50%	37.80%	1.12

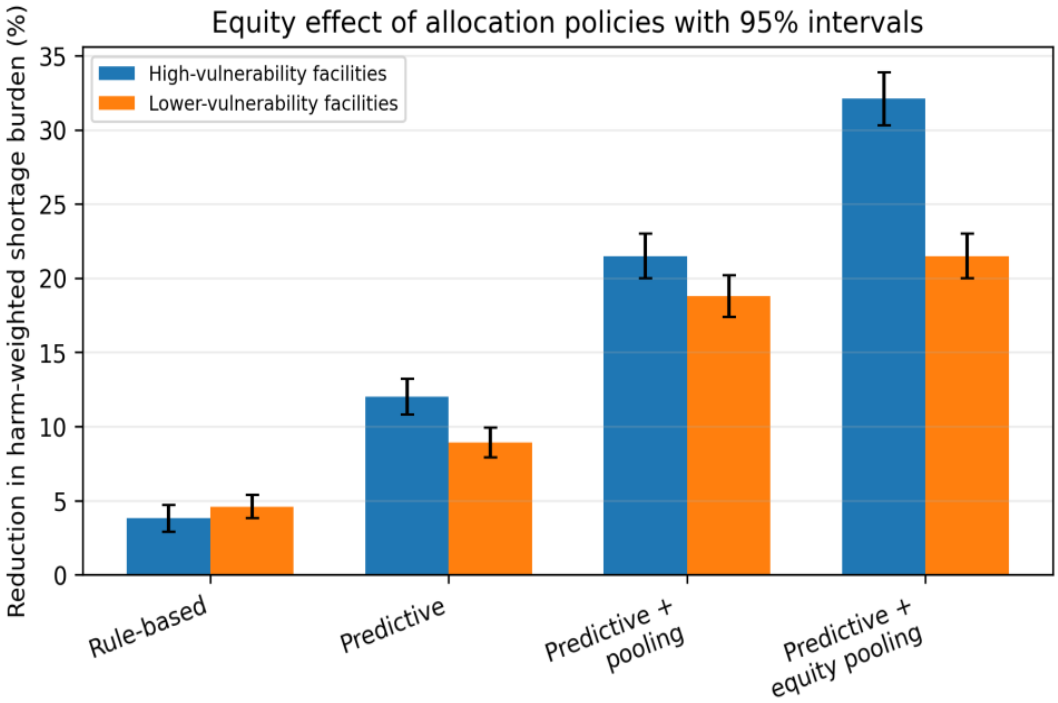


Figure 5. Harm-weighted shortage-burden reduction by facility vulnerability group. Error bars show 95% simulation intervals.

CASE INTERPRETATION

The simulated network begins the year with uneven inventory resilience. Critical access hospitals and rural emergency hospitals face high demand for emergency and injectable medicines, but have longer replenishment routes. FQHCs and safety-net outpatient pharmacies have lower acute-care demand but high exposure to disruptions in chronic disease medication access because their patients may have few alternative access points. The tribal health clinic receives the highest rural weight and therefore becomes more protected under equity-weighted allocation.

Under the reactive policy, shortage signals are recognised late, and facilities submit emergency orders after stockouts have already occurred. Rule-based inventory triggers reduce some local stockout severity but still miss upstream deterioration. Predictive analytics improves preparedness by moving high-risk medicines into earlier review. Regional pooling then converts information into action by redistributing feasible surplus. Equity-weighted pooling prevents the predictable injustice of allocating scarce medicine only by institutional size, historical purchasing volume or existing market power.

The results, therefore, support a narrow but important conclusion: synthetic simulation demonstrates that the proposed decision logic is operationalisable and directionally plausible. It does not prove real-world effectiveness. Real-world effectiveness requires facility-level data, external validation, clinical oversight and implementation evaluation.

DISCUSSION

The findings support an anticipatory view of medication-shortage resilience. Traditional shortage management asks what should be done after a medicine becomes unavailable. The framework developed here asks which medicines are likely to become locally unavailable, which facilities are most vulnerable, which patients may be harmed, and which intervention should be triggered before the shortage becomes clinically disruptive.

The first major implication is that predictive accuracy alone is insufficient. This is methodologically incomplete because predictive accuracy does not establish operational value. A model with good discrimination but no link to procurement, substitution, regional transfer or clinical governance may perform well in a technical notebook and fail in practice. For this reason, the revised manuscript reports operational outcomes alongside model performance.

The second implication is that regional coordination increases the value of prediction. Many rural and underserved facilities cannot self-insure with large inventories because storage, budget, expiry, cold-chain, and supplier allocation constraints are binding. A regional network allows surplus stock to be matched with predicted need before crisis conditions emerge.

The third implication is that equity must be designed into the decision rule. A model trained only on historical purchasing data risks reproducing historical under-service. Smaller facilities may appear less important because they purchased fewer units in the past, even when the unmet need is high. The equity-weighted framework corrects this by considering vulnerability, rurality, clinical criticality and lack of alternatives.

POLICY AND MANAGERIAL IMPLICATIONS

1. National shortage reporting should be complemented by local operational surveillance. FDA and ASHP signals are necessary but insufficient for facility-level decision-making, as fill rates, lead times, order cancellations, and days on hand are often more actionable early-warning indicators.
2. Rural facilities should not be expected to build analytics infrastructure alone. State health departments, hospital associations, academic health systems, group purchasing organisations and health information exchanges could host shared shortage-risk dashboards for resource-constrained providers.
3. Procurement policy should reward reliability as well as price. Contracting systems that emphasise price alone can create fragile supply conditions and shift the burden of resilience to pharmacists and clinicians.
4. Regional mutual-aid agreements should be formalised before shortages occur. Ad hoc borrowing during a crisis is inefficient and ethically inconsistent. Predefined rules are needed for stock-sharing, documentation, transport, liability and reimbursement.
5. Allocation governance should include equity metrics. Historical volume, institutional size and purchasing power should not dominate access when clinically vulnerable and geographically isolated populations are at risk.

LIMITATIONS

The central limitation is that the evaluation is synthetic. The simulation is transparent and useful for testing decision logic, but it is not a substitute for real-world evaluation. Actual hospital data may reveal different demand distributions, supplier behaviours, transfer constraints, payer barriers, substitution practices, clinician preferences or patient-level consequences.

The synthetic parameters may introduce structural bias. The model also simplifies legal, contractual, 340B, payer-related, cold-chain, expiry and liability constraints that may restrict interfacility stock transfer. The

treatment-delay index is a proxy outcome; patient-level harm should be measured directly where ethically and legally feasible.

The framework does not solve upstream manufacturing fragility. Predictive analytics cannot manufacture sterile injectables, repair weak generic-drug economics or replace federal coordination. It can reduce local harm while broader manufacturing, regulatory and procurement reforms proceed.

There is also a risk of algorithmic bias. Historical purchasing data may encode under-access. Future empirical deployment must therefore include calibration review, subgroup performance monitoring, clinician oversight, fairness auditing and transparent appeal mechanisms.

FUTURE EMPIRICAL STUDY DESIGN

The next stage should be a multi-site empirical study. A stepped-wedge design would be appropriate because all participating facilities could eventually receive the intervention, while staggered rollout permits comparison. Data should be collected from pharmacy inventory systems, purchasing records, EHR medication administration records, wholesaler fill-rate feeds, national shortage databases and regional transfer logs.

The primary endpoint should be local stockout days for high-criticality medicines per facility-month. Secondary endpoints should include shortage units, treatment-delay index, emergency procurement cost, substitution frequency, pharmacist workload, clinician satisfaction, patient travel burden and equity-weighted disruption burden.

A rigorous empirical analysis should include facility fixed effects, drug fixed effects, time effects and interaction terms for rurality, clinical criticality and baseline vulnerability. Model performance should be reported alongside decision impact because prediction accuracy is not equivalent to healthcare value.

CONCLUSION

Medication shortages are persistent public health and health equity challenges in the United States. Rural and underserved facilities are particularly vulnerable because they operate with constrained inventory, limited supplier leverage, thinner staffing, longer replenishment routes and patient populations with fewer alternatives. This revised paper proposes a predictive analytics framework that links shortage-risk prediction to facility vulnerability, inventory visibility, regional pooling and equity-sensitive allocation governance. The synthetic simulation-based evaluation shows that the strongest policy, predictive analytics with equity-weighted pooling, reduces shortage-unit severity, treatment-delay burden, emergency orders and disruption cost relative to reactive management. The conclusion is deliberately bounded: the framework is operationally specified, and the simulation results are directionally plausible, but real-world effectiveness must be established through multi-site empirical testing using hospital pharmacy, wholesaler, and EHR data.

Data Availability Statement

The case-study environment used in this manuscript is synthetic and was generated to demonstrate the proposed framework. No patient-level data, real facility inventory records or protected health information were used. For journal submission, the synthetic data-generation code, parameter table and simulation outputs should be deposited in a repository. If the study is repeated with real facility data, data sharing will depend on institutional review board approval, data-use agreements and privacy requirements.

Funding Statement

No external funding is declared for this manuscript draft.

Conflict of Interest Statement

The author declares no conflict of interest in relation to this conceptual and simulation-based manuscript draft.

REFERENCES

1. U.S. Food and Drug Administration. Drug Shortages: Drug Safety and Availability. Updated October 23, 2025. Accessed June 15, 2026.
2. American Society of Health-System Pharmacists. Drug Shortages Statistics. Accessed June 15, 2026.
3. U.S. Government Accountability Office. Drug Shortages: HHS Should Implement a Mechanism to Coordinate Its Activities. GAO-25-107110. Published April 9, 2025.
4. National Academies of Sciences, Engineering, and Medicine. Building Resilience into the Nation's Medical Product Supply Chains. Washington, DC: The National Academies Press; 2022.
5. U.S. Department of Health and Human Services. 2025-2028 Draft Action Plan for Addressing Shortages of Medical Products and Critical Foods and Strengthening the Resilience of Medical Product and Critical Food Supply Chains. Published January 2025.
6. McGeeney JD, et al. Analysis of Drug Shortages, 2018-2023. National Academies/NCBI Bookshelf; 2025.
7. Tucker EL, Daskin MS. Pharmaceutical supply chain reliability and effects on drug shortages. Computers & Industrial Engineering. 2022;169:108258.
8. Sabogal de la Pava ML, Tucker EL. Effects of geopolitical strain on global pharmaceutical supply chain design and drug shortages. European Journal of Operational Research. 2025;327(2):641-654.
9. Kosmas D, Ergun O. When should the FDA inspect pharmaceutical manufacturing facilities to better mitigate drug shortages? Working paper; 2023.
10. Cui M, Jiang Y, Zhou D, Qian C, Zhang Y, Wang Q. ShortageSim: Simulating drug shortages under information asymmetry. arXiv preprint; 2025.
11. Arji G, Ahmadi H, Avazpoor J, Hemmat M. Identifying resilience strategies for disruption management in healthcare supply chains: a systematic review. BMC Health Services Research. 2023.
12. Feng Q, Du M, Zou N, Hu X. Fair machine learning in healthcare: a review. arXiv preprint; 2022.
13. Esnaashari S, Bright J, Francis J, Hashem Y, Straub V, Morgan D. Approaches to the algorithmic allocation of public resources: a cross-disciplinary review. arXiv preprint; 2023.
14. Commonwealth Fund. The State of Rural Primary Care in the United States. Published November 17, 2025.
15. Chartis Centre for Rural Health. 2025 Rural Health State of the State. Published February 10, 2025.
16. Sheps Centre for Health Services Research. Rural Hospital Closures. University of North Carolina. Accessed June 15, 2026.
17. Aishat Olutosin Olukotun (2026). Integrating Demand Forecasting, Real-Time Inventory Visibility, And Equitable Allocation Models to Mitigate Medication Shortages in U.S. Healthcare Supply Chains. , 13(15), https://doi.org/10.51244/IJRSI.2026.1315PH00149
18. Silver EA, Pyke DF, Thomas DJ. Inventory and Production Management in Supply Chains. 4th ed. CRC Press; 2016.
19. Nahmias S, Olsen TL. Production and Operations Analysis. 7th ed. Waveland Press; 2015.
20. Alasi Bashirat Ololade. Improving Diagnostic Turnaround Time and Analytical Reliability in Clinical Biochemistry Laboratories: A Quality Management Framework. (2026). International Journal of Latest Technology in Engineering Management & Applied Science, 15(6),

APPENDIX A. AMENDMENT MATRIX

Appendix Table A1. How the major review concerns were addressed in the revised manuscript.

Critical review issue	Amendment made
Overstatement of validation	Replaced validation language with synthetic simulation-based evaluation and explicitly bounded the claims.
Underdeveloped mathematical model	Added full notation and ordered equations for forecasting, visibility, risk, allocation and evaluation.
Unclear simulation design	Added Algorithm 1 and a parameter-assumption table explaining synthetic data generation.

Prediction-performance confusion	Separated model-performance metrics from policy-outcome comparisons.
Weak equity analysis	Added harm-weighted equity outcomes and vulnerability-group comparison.
Figures lacked uncertainty	Replaced simple bars with plots showing 95% simulation intervals.
Informal academic tone	Removed conversational or dismissive language and replaced it with disciplined scholarly phrasing.
Limitations too narrow	Expanded limitations to include synthetic-parameter bias, legal constraints, cold-chain, 340B/payer barriers, patient-level outcomes and algorithmic bias.
Future empirical design is too broad.	Specified a primary endpoint and secondary endpoints for a future stepped-wedge study.