Integrating Demand Forecasting, Real-Time Inventory Visibility, And Equitable Allocation Models to Mitigate Medication Shortages in U.S. Healthcare Supply Chains

Aishat Olutosin Olukotun

University of Ibadan

DOI: https://doi.org/10.51244/IJRSI.2026.1315PH00149

Received: 22 July 2026; Accepted: 28 July 2026; Published: 05 August 2026

ABSTRACT

Medication shortages in the United States are not just occasional supply failures; they are ongoing disruptions to the health system caused by poor demand visibility, limited upstream transparency, fragile manufacturing economics, and fragmented coordination. Public data illustrates the problem's extent. ASHP reported 223 active drug shortages recently, down from the 323 at the peak in early 2024 but still rising quarter by quarter. The FDA's 2024 report to Congress revealed 15 new shortages in 2024 and 113 ongoing shortages as of December 2024, 2024. Independent analysis of data from 2018-2023 shows shortages increased from 313 in January 2020 to 350 in December 2023, with injectable medicines being most affected. This article proposes a decision-support framework that combines three often-separated components: probabilistic demand forecasting, real-time inventory visibility, and equity-based allocation of scarce medicines. Its contribution focuses on methodology rather than purely on empirical evidence, outlining the minimum data requirements, mathematical formulation, optimisation logic, governance conditions, and evaluation metrics needed to shift shortage management from reactive substitution to proactive resilience. The framework aims to serve rural hospitals, safety-net providers, pharmacies, and high-acuity services where drug shortages can quickly cause treatment delays, urgent procurement, and unfair access. The paper concludes that an effective shortage mitigation system must integrate forecast calibration, facility-level risk scoring, and transparent allocation rules; otherwise, it is merely a dashboard with enhanced visuals.

Keywords: medication shortages; pharmaceutical supply chain; demand forecasting; inventory visibility; equitable allocation; rural hospitals; safety-net providers; sterile injectables; healthcare analytics; supply chain resilience.

INTRODUCTION

Drug shortages have become a persistent failure in the U.S. healthcare supply chain. They affect hospitals, retail and speciality pharmacies, oncology centres, emergency departments, nephrology services, urology practices and safety-net providers. Their effect is not administrative inconvenience. It is delayed treatment, therapeutic substitution, cancelled procedures, avoidable labour costs, emergency procurement, and sometimes rationing decisions made under clinical pressure. A competent manuscript must therefore stop describing shortages as mere purchasing disturbances and treat them as a coupled forecasting, visibility and allocation problem. Recent public evidence confirms both the persistence and complexity of the problem. ASHP and the University of Utah Drug Information Service reported 223 active shortages in a recent public snapshot, lower than the record 323 active shortages reported in the first quarter of 2024 but still substantial. FDA states that shortages arise from manufacturing and quality problems, delays and discontinuations. At the same time, GAO concluded in 2025 that HHS should implement a formal mechanism to coordinate drug-shortage activities and collaborate with federal stakeholders [1-5]. The implication is obvious: no single hospital, wholesaler or manufacturer can solve this alone. However, a facility-level decision system can reduce preventable harm if it is built on real demand signals, real inventory status and defensible allocation logic.

The central thesis of this article is simple: demand forecasting, inventory visibility, and equitable allocation must be integrated into a single operational architecture. Forecasting without inventory visibility is academic ornamentation. Inventory visibility without allocation logic merely informs users that scarcity exists. Allocation guidance without real-time data degenerates into improvised rationing. The contribution of this manuscript is to specify a rigorous framework that binds all three components into a model that can be tested, audited and deployed.

Empirical Motivation and Public Data Anchors

This article does not invent private hospital results. That would be poor science. Instead, it uses public shortage indicators to motivate the modelling framework and then defines the data needed for empirical implementation. Table 1 consolidates the public evidence used to justify the proposed model.

Table 1. Public data anchors motivating the integrated shortage-mitigation framework

Source	Data element	Reported value	Analytical implication
ASHP active-shortage statistics	Active shortages	223 active shortages; lower than the Q1 2024 peak of 323 but trending upward in the latest public snapshot	Persistent national shortage burden; supports early-warning and allocation logic [1].
ASHP 2024 national shortage report	End-of-year active shortages	271 active shortages at the end of 2024; nearly half began in 2022 or earlier; Hurricane Helene was associated with 12 new product shortages of critical fluids	Shortages are durable and disruption-sensitive, not merely short-lived stock anomalies [2].
FDA Drug Shortages page	Root causes	FDA identifies manufacturing and quality problems, delays and discontinuations as the causes	Forecasting must include upstream production and quality-risk signals, not demand alone [3].
FDA Report to Congress CY2024	New and ongoing shortages	15 new drug shortages in CY2024; 113 ongoing FDA-tracked shortages as of 31 Dec 2024	New-shortage counts alone understate the burden because old shortages persist [4].
NCBI/IQVIA-based analysis, 2018-2023	Ongoing shortage events	Ongoing shortage events increased from 313 in Jan 2020 to 350 in Dec 2023; injectable events rose from 208 to 251	Injectables are a high-priority modelling category for hospitals [8].
GAO 2025	Governance and coordination	GAO recommended a formal HHS mechanism for coordinating drug-shortage activities across stakeholders	Analytics platforms require governance and cross-sector coordination, not isolated dashboards [5].

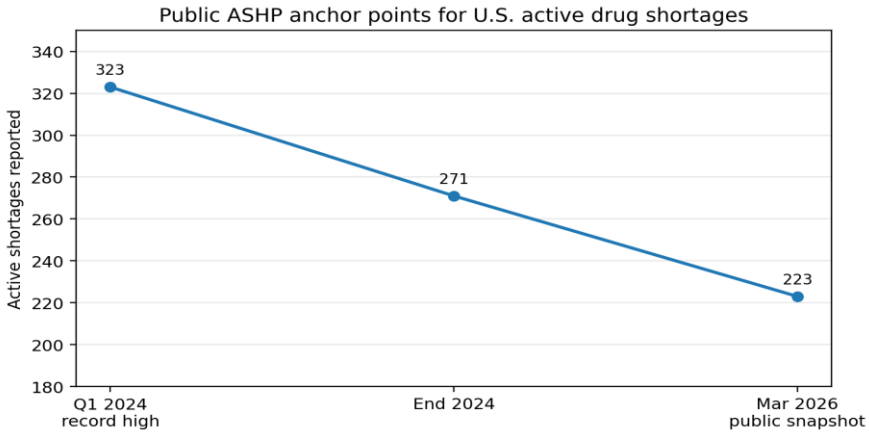

Figure 1. Public ASHP active-shortage anchor points. Note: these are public anchor points from ASHP summaries, not a continuous quarterly time series.

Figure 1 should be read carefully. It does not claim that shortages declined monotonically or that March 2026 is harmless. It shows that even after the Q1 2024 peak, active shortages remained at a level requiring systematic mitigation. The intellectually lazy interpretation is to celebrate a fall from 323 to 223. The correct interpretation is that a system with hundreds of active shortages still needs predictive and allocative infrastructure.

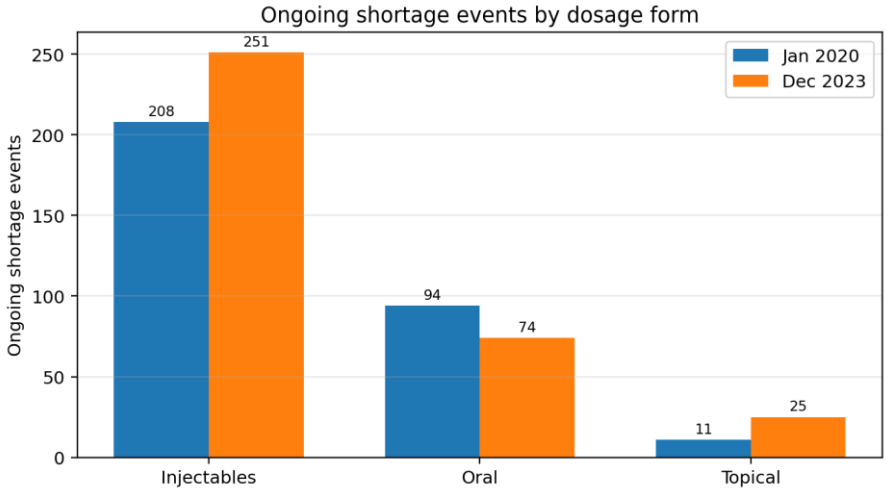

Figure 2. Ongoing shortage events by dosage form, January 2020 versus December 2023, based on the NCBI Bookshelf analysis of FDA shortage data.

Figure 2 highlights why hospital-oriented models must treat injectables as a priority class. Injectable shortage events increased from 208 to 251 between January 2020 and December 2023, while oral shortage events declined from 94 to 74, and topical events increased from 11 to 25. A forecasting model that treats all dosage forms identically is technically convenient but clinically naive.

Data Requirements for an Empirical Platform

A publishable empirical extension requires transaction-level and facility-level data. The platform should ingest public shortage warnings, internal utilisation data, ERP stock data, supplier lead-time information and facility vulnerability indicators. The minimum viable data dictionary is shown in Table 2.

Table 2. Minimum viable data dictionary for the proposed shortage-mitigation platform.

Variable	Definition	Likely source	Use in the model
facility_id	Facility identifier	Hospital, pharmacy or clinic master file	Needed for allocation, equity auditing and regional analysis.
drug_id / NDC / generic_name	Medicine identifier	ERP, pharmacy inventory, FDA/ASHP mapping	Links shortage signals to actual products purchased or dispensed.
daily_utilization	Units dispensed or administered per day	EHR, pharmacy dispensing, charge capture	Primary dependent variable for forecasting.
stock_on_hand	Available usable inventory	ERP, pharmacy inventory system, automated dispensing cabinet	Core visibility input; must exclude quarantined and expired stock.
inbound_quantity and ETA	Confirmed supply on order and delivery date	Wholesaler portal, purchase orders, ERP	Separates real shortage risk from temporary local low-stock events.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue XV July 2026 | Issue on Public Health

lead_time_days	Supplier or wholesaler replenishment delay	ERP purchase-order history	Controls reorder point and safety-stock calculations.
substitution_index	Clinical substitutability score	Pharmacy and therapeutics committee, clinical guidelines	Prevents a purely economic model from treating critical medicines as interchangeable.
clinical_criticality	Acuity and severity score	Clinical governance, formulary classification	Weight shortage causes harm due to clinical consequences.
equity_vulnerability	Rurality, safety-net share, payer mix, area deprivation index	Facility attributes, CMS/HHS datasets	Prevents allocation rules from privileging already-resourced facilities.
stockout_event	Binary/continuous shortage outcome	Inventory and dispensing logs	Used for model validation and KPI reporting.

Integrated Decision Architecture

The proposed architecture has four layers: data ingestion, forecasting, shortage-risk scoring and equity-weighted allocation. The sequence matters. If the model begins with allocation before estimating demand uncertainty and inventory state, the formulation is mathematically backwards. Figure 3 shows the required architecture.

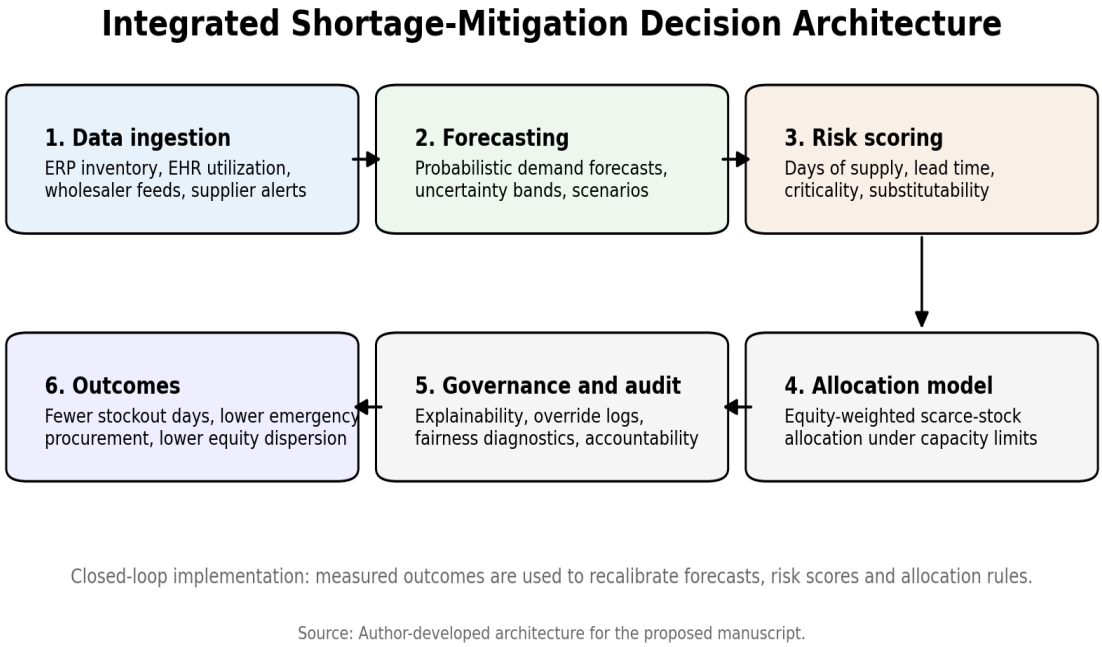

Figure 3. Integrated shortage-mitigation decision architecture linking data ingestion, forecasting, risk scoring, allocation, governance and evaluation

The architecture is deliberately modular. The forecasting layer can use statistical time-series models, gradient boosting, Bayesian hierarchical models or recurrent neural networks, provided calibration is tested. The allocation layer should not be a black box. It must expose the weights assigned to clinical criticality, facility vulnerability, lead time and substitutability. In healthcare supply chains, explainability is not decoration; it is governance.

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue XV July 2026 | Issue on Public Health

Ordered Mathematical Formulation

The following formulation gives the equations in the correct order. Sets, parameters and decision variables are defined first, followed by demand forecasting, uncertainty estimation, inventory visibility, shortage-risk scoring, allocation optimisation, and evaluation metrics. This order is not aesthetic; it reflects logical dependencies. Allocation decisions require demand estimates and available inventory. Risk scores require days of supply and criticality. Evaluation metrics require observed outcomes.

Sets, indices and notation

Table 3. Notation used in the mathematical formulation.

Symbol	Meaning
F	Set of facilities indexed by f, including hospitals, pharmacies and clinics.
D	Set of medicines or drug products indexed by d.
T	Set of decision periods indexed by t.
S	Set of demand or disruption scenarios indexed by s.
D_{fdt}	Observed demand or utilisation for facility f, drug d and period t.
$Dhat_{fdt}$	Forecast demand for facility f, drug d and period t.
I_{fdt}	Usable inventory available at facility f for drug d at period t.
Q_{fdt}	Confirmed inbound order quantity for facility f and drug d at period t.
x_{fdst}	Allocation quantity sent to facility f for drug d under scenario s.
u_{fdst}	Unserved demand or shortage quantity.
V_f	Facility vulnerability score, reflecting rurality, safety-net burden or deprivation.
C_d	Clinical criticality score for drug D.
A_d	Available central supply of drug D for allocation.

Figure 4. Correct sequence of equations in the proposed model.

Demand forecasting and uncertainty

$$Dhat_{\{f,d,t+h\}} = E[D_{\{f,d,t+h\}} \mid X_{\{f,d,t\}}] \quad (1)$$

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue XV July 2026 | Issue on Public Health

Equation (1) defines the h-step-ahead conditional demand forecast using covariates X such as utilisation trend, seasonality, procedure volume, prescribing pattern, public shortage alerts and local epidemiological signals.

$$\text{sigmahat}_{\{f,d,t+h\}} = \text{sqrt}(E[(D_{\{f,d,t+h\}} - \text{Dhat}_{\{f,d,t+h\}})^2]) \quad (2)$$

Equation (2) estimates forecast uncertainty. A shortage platform that reports only point forecasts is not fit for decision-making under scarcity.

Inventory visibility and days of supply

$$I_{\{f,d,t+1\}} = I_{\{f,d,t\}} + Q_{\{f,d,t\}} + x_{\{f,d,t\}} - D_{\{f,d,t\}} - E_{\{f,d,t\}} \quad (3)$$

Equation (3) is the facility-level inventory balance. E represents expiry, quarantine, recall loss or unusable inventory. This is where real-time visibility enters the model.

$$\text{DOS}_{\{f,d,t\}} = I_{\{f,d,t\}} / \max(\text{epsilon}, \text{Dhat}_{\{f,d,t+1\}}) \quad (4)$$

Equation (4) converts stock into days of supply. The small epsilon prevents division by zero for low-utilisation medicines.

Shortage-risk score

$$SR_{\{f,d,t\}} = w1*P(D_{\{f,d,t+1\}} > I_{\{f,d,t\}}+Q_{\{f,d,t\})} + w2*L_{\{d,t\}} + w3*C_d + w4*(1-SUB_d) + w5*V_f \quad (5)$$

Equation (5) produces a shortage-risk score from probabilistic demand exceedance, lead time L, clinical criticality C, lack of substitutability SUB and facility vulnerability V. The weights w must be specified, governed and tested, not hidden in a spreadsheet.

Equity-weighted allocation optimization

$$\text{minimize } Z = \text{sum}_s\, pi_s\, [\, \text{sum}_f\, \text{sum}_d\, H_{\{f,d\}}\, u_{\{f,d,s\}} + \text{lambda}\, \text{sum}_d\, ED_{\{d,s\}} + \text{mu}\, \text{sum}_f\, \text{sum}_d\, c_{\{f,d\}}\, x_{\{f,d,s\}}\,] \quad (6)$$

Equation (6) minimises the expected harm from shortages. The first term penalises unmet, clinically weighted demand; the second penalises inequity in dispersion; and the third accounts for distribution costs.

$$\text{sum\_f } x_{\{f,d,s\}} <= A_{\{d,s\}} \text{ for all d, sq} \quad (7)$$

Equation (7) enforces central supply availability. No allocation model can distribute inventory that does not exist; this is where naive policy prose becomes arithmetic.

$$I_{\{f,d,t\}} + Q_{\{f,d,t\}} + x_{\{f,d,s\}} + u_{\{f,d,s\}} >= D_{\{f,d,s\}} \text{ for all f, d, s} \quad (8).$$

Equation (8) links allocation and unserved demand. If available stock and allocation are insufficient, the shortage variable u becomes positive.

$$0 <= x_{\{f,d,s\}} <= CAP_{\{f,d\}} \text{ for all f, d, s} \quad (9)$$

Equation (9) prevents infeasible allocations that exceed storage limits, handling capacity or clinically usable inventory levels at a facility.

$$r_{\{f,d,s\}} = (I_{\{f,d,t\}}+Q_{\{f,d,t\}}+x_{\{f,d,s\}}) / \max(\text{epsilon}, D_{\{f,d,s\}}) \quad (10)$$

Equation (10) defines the demand-coverage ratio used for fairness evaluation and allocation audit.

$$r_{\{f,d,s\}} >= \text{rho}_{\min} * V_f * C_d \text{ for priority facilities and critical drugs} \quad (11)$$

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue XV July 2026 | Issue on Public Health

Equation (11) imposes an equity-sensitive floor, ensuring that vulnerable facilities and critical drugs receive at least the minimum coverage. The parameter rho_{min} should be determined by governance policy, not by the analyst alone.

Evaluation metrics

$$ED_{\{d,s\}} = \max_f(r_{\{f,d,s\}}) - \min_f(r_{\{f,d,s\}}) \quad (12)$$

Equation (12) measures allocation dispersion. A high value signals unequal coverage across facilities.

$$\text{StockoutRate} = \text{sum}_{\{f,d,t\}} 1(I_{\{f,d,t\}}=0 \text{ and } D_{\{f,d,t\}}>0) / |F||D||T| \quad (13)$$

Equation (13) measures the frequency of observed stockout exposure.

$$MAPE_d = (1/n)\ \text{sum}_t\ |D_{\{d,t\}} - Dhat_{\{d,t\}}| / \max(\text{epsilon}, D_{\{d,t\}}) \quad (14)$$

Equation (14) evaluates forecast error. Calibration metrics should complement it because shortages are tail-risk events.

$$\text{Emergency Cost Reduction} = (C_{base} - C_{model}) / C_{base} \quad (15)$$

Equation (15) evaluates whether the platform reduces emergency procurement cost relative to a baseline policy.

Model Implementation and Pilot Design

The empirical test should be staged in two phases. First, a retrospective study should reconstruct historical shortage episodes using hospital, pharmacy or wholesaler records. Second, a prospective pilot should run the platform in shadow mode before any automated allocation authority is granted. Shadow mode is essential because clinicians and pharmacists must see whether the model produces sensible alerts before it changes real supply decisions.

Table 4. Empirical validation pathway for the proposed platform.

Pilot phase	Data requirement	Primary evaluation endpoints
Phase 1: Retrospective validation	12-36 months of utilisation, inventory and purchase-order history; public shortage dates	Forecast accuracy, predicted stockout probability, lead-time accuracy, and shortage-risk calibration.
Phase 2: Shadow-mode pilot	Live inventory feeds and demand forecasts; no automatic reallocation	Alert precision, false-alarm burden, pharmacist usability, and dashboard latency.
Phase 3: Assisted allocation	Governance-approved allocation recommendations; human override allowed	Stockout reduction, emergency procurement cost, treatment delay, and equity dispersion.
Phase 4: Scaled deployment	Multi-facility network with dashboard, audit trail and periodic model retraining	System resilience, cross-facility fairness, supplier response, and operational sustainability.

Dashboard and Governance Requirements

A dashboard is useful only if it changes decisions. The dashboard should display forecasted demand, usable stock, days of supply, inbound orders, substitution options, shortage risk score, allocation recommendation, and the reason code for each recommendation. Each override should be logged with the user role, timestamp, reason

and expected clinical consequence. This is not bureaucratic excess. It is what prevents algorithmic allocation from becoming unaccountable rationing.

- Data quality controls: missing-stock flags, stale inventory warnings, duplicate NDC mapping checks and expiry-date exclusions.
- Clinical controls: formulary substitution rules, drug criticality classification and service-line prioritisation.
- Equity controls: rurality and safety-net indicators, deprivation indices and monitoring of allocation dispersion.
- Operational controls: lead-time uncertainty, minimum order quantities, storage constraints, cold-chain requirements and supplier reliability.
- Audit controls: model versioning, override logs, performance drift alerts and periodic stakeholder review.

DISCUSSION

The framework advances the medication-shortage literature in three ways. First, it reframes shortage mitigation as a closed-loop decision problem rather than a forecasting exercise. Second, it makes equity operational by embedding vulnerability and criticality within allocation constraints rather than treating them as rhetorical commitments. Third, it proposes evaluation metrics that journals and policymakers can interrogate: stockout rate, emergency procurement cost, treatment delay, forecast calibration and equity dispersion. The model is also intentionally modest. It does not claim that analytics can manufacture unavailable medicines. It does not pretend that dashboards can repair fragile generic-drug economics. What it can do is reduce avoidable harm inside the healthcare delivery system by anticipating local shortage risk and distributing scarce stock more transparently. That is a defensible contribution. Do not oversell it.

Policy and Managerial Implications

For policymakers, the framework supports calls for stronger supply-chain visibility and formal coordination. GAO's recommendation for a formal HHS coordination mechanism is consistent with the need for shared metrics and cross-stakeholder data exchange. For hospital leaders, the model justifies investment in ERP integration, pharmacy inventory data quality and analytics governance. For pharmacists, it provides a structured method for prioritising scarce stock without relying only on manual workarounds. For rural and safety-net providers, the equity floor in Equation (11) addresses a core concern: allocation systems should not simply reward facilities with more purchasing power, better analysts or larger buffers.

Limitations and Future Research

The main limitation is data availability. Many facilities lack clean, real-time inventory feeds, and wholesaler data may not be consistently available. Demand forecasts may also be biased if historical utilisation reflects prior access constraints rather than true clinical need. Equity scores require careful governance because poorly chosen proxies can create new unfairness. Future research should validate the model using multi-site data, compare forecasting algorithms, estimate the causal effect of allocation recommendations on stockouts and test whether equity floors improve access without producing unacceptable inefficiency.

CONCLUSION

Medication shortages are a public-health supply-chain problem that requires more than reactive purchasing. A high-impact response must integrate probabilistic demand forecasting, real-time inventory visibility and equitable allocation. This article provides a data-informed, mathematically structured framework for doing so. The next step is empirical validation using real facility-level data. Until then, the honest claim is not that the model has solved shortages, but that it specifies the decision architecture required to reduce avoidable harm in a system where scarcity is recurrent.

Operationalisation of the Proposed Platform

A credible shortage-mitigation platform cannot be built as a collection of disconnected reports. It must operate as a closed-loop decision system in which data ingestion, forecasting, risk classification, allocation, feedback and governance are linked through repeatable operating rules. The practical implementation should therefore begin with a limited set of high-value drugs rather than the entire formulary. A sensible pilot would prioritise sterile injectables, oncology medicines, renal-support drugs, antimicrobials, and other agents for which substitution during shortages is clinically difficult or operationally expensive. This is not administrative fussiness; it is the discipline that separates an implementable decision system from a fashionable analytics dashboard.

The first operational step is master data harmonisation. National Drug Codes, facility formularies, package sizes, therapeutic classes, supplier identifiers and substitute products must be mapped into a common data model. Without that mapping, an apparent inventory deficit may be a coding problem, and an apparent surplus may be unusable because it is in the wrong package size, dose form or care setting. The platform should therefore maintain a controlled drug dictionary that links product identifiers to clinical criticality, substitution rules, storage conditions, expiry constraints and shortage status.

The second operational step is data-latency control. Real-time visibility is only meaningful when the age of each data feed is known. A facility whose inventory file was refreshed five minutes ago should not be treated the same as a facility whose record is two days old. The platform should therefore attach timestamp metadata to each inventory, order, dispensing and receipt record. Stale records should be flagged, excluded from automated allocation where necessary, and routed for human verification. Peer reviewers will immediately penalise a paper that claims real-time inventory visibility but fails to define data freshness.

Data integration and validation workflow

The proposed data pipeline should follow five validation gates. Gate 1 confirms that the facility, drug and supplier identifiers are valid. Gate 2 checks inventory arithmetic by comparing beginning stock, receipts, dispensing, transfers, expiry, and ending stock, Gate 3 tests whether demand observations are plausible relative to recent utilisation and service-line activity. Gate 4 detects missing or stale feeds. Gate 5 compares internal shortage-risk predictions with external shortage signals from public or supplier-facing sources. These gates are not optional. They are the minimum hygiene requirements before optimisation can influence patient-facing allocation. For each facility-drug pair, the platform should compute a data-reliability score. This score should not be confused with shortage risk. A facility may have low inventory risk but poor data reliability, or high shortage risk with excellent data quality. The difference matters because the recommended managerial action differs. Low reliability calls for data correction; high shortage risk calls for procurement, substitution, redistribution or allocation intervention. Combining both into one score would be intellectually lazy and operationally dangerous. A suggested reliability score may combine feed freshness, missingness, inventory-balance consistency, duplicate-code frequency and exception-rate history. In a future empirical study, this score can be used as a covariate in the shortage-risk model or as a constraint in the allocation model. Facilities with unreliable data should not be denied support, but their recommendations should trigger human review. This is how analytics becomes responsible governance rather than blind automation.

Forecasting design for intermittent and crisis-sensitive demand

Medication demand is not always smooth. Some drugs show regular utilisation patterns, while others are intermittent, seasonal, protocol-driven or event-sensitive. A single forecasting method is therefore inadequate. The platform should allow model selection by demand class. Stable high-volume products may be forecast using classical time-series or regularised regression models. Intermittent products may require Croston-type methods, Bayesian shrinkage or hierarchical pooling across facilities. Crisis-sensitive products may require scenario overlays that incorporate epidemiological signals, procedure schedules, recall notices, supplier disruptions and policy changes. Forecast uncertainty is as important as the point forecast. The purpose of the forecast is not merely to estimate expected demand but to quantify the probability that demand will exceed usable stock before replenishment arrives. This is a tail-risk problem. A model that performs well on average error but fails during

ISSN No. 2321-2705 | DOI: 10.51244/IJRSI |Volume XIII Issue XV July 2026 | Issue on Public Health

rare demand surges is not fit for shortage mitigation. The evaluation section should therefore report calibration, prediction-interval coverage and shortage-event recall, not only mean absolute percentage error. This is elementary, but many papers in applied healthcare analytics still get it wrong.

The forecasting layer should also support hierarchical reconciliation. Drug demand at a facility is nested within service lines, regional referral patterns, supplier contracts and national shortage conditions. Reconciled forecasts can prevent inconsistent recommendations, such as predicting a regional decline even though all major facilities in that region are forecast to increase utilisation. The reconciled structure also helps smaller rural facilities by borrowing information from clinically similar institutions without erasing local differences.

Equity-Aware Allocation in Practice

Equitable allocation is not the same as equal allocation. Equal allocation distributes scarce medicines in the same proportion or quantity across facilities, regardless of clinical need, substitution difficulty or structural disadvantage. Equity-aware allocation recognises that facilities differ in baseline access, patient vulnerability, travel distance, payer mix, service-line dependency and substitute availability. A safety-net hospital serving patients with limited transportation options is not operationally equivalent to a large urban system with multiple supplier relationships and stronger purchasing leverage. The model must encode this difference explicitly and transparently. The equity weight in the allocation formulation should be constructed from interpretable components. Examples include rurality, safety-net status, proportion of Medicaid or uninsured patients, local deprivation, distance to alternative care, historical fill-rate disadvantage, therapeutic criticality and absence of clinically acceptable substitutes. However, these weights must be auditable. They should not be hidden inside a black-box score that stakeholders cannot interrogate. In practice, the platform should display both the final equity weight and the factors contributing to it.

The allocation model should also distinguish between preventive and emergency allocation. Preventive allocation moves inventory before a stockout occurs, based on forecasted risk. Emergency allocation responds to confirmed exposure to shortages. The ethical logic differs. Preventive allocation can tolerate more uncertainty because the intervention is designed to reduce risk; emergency allocation requires stronger evidence because scarce stock is being redirected from one facility to another. A high-impact paper should state this distinction clearly because it has direct governance implications.

Allocation audit and fairness diagnostics

After each allocation cycle, the platform should generate a fairness audit. At a minimum, the audit should compare coverage ratios across facility classes, drug criticality classes and vulnerability quartiles. It should report whether rural and safety-net facilities experienced lower residual shortage exposure than under a baseline first-come-first-served or proportional allocation rule. It should also identify facilities that repeatedly receive lower coverage after adjustment for demand, clinical criticality, and storage capacity. Repeated disadvantage is a signal of structural inequity, not merely a statistical inconvenience.

Fairness diagnostics should include both group-level and facility-level measures. Group-level measures reveal whether vulnerable categories are systematically underserved. Facility-level measures reveal whether individual hospitals or pharmacies are repeatedly penalised by data quality issues, supplier access, geographic isolation, or small order volumes. In a journal submission, this distinction is important because a model can look fair on average while still harming particular facilities. A serious reviewer will notice this immediately. The audit process should also record override decisions. Human override is not a weakness in clinical supply chains; it is often necessary. The weakness is allowing override decisions to occur without documentation. Every override should record the reason, the actor, the affected product, the affected facility, the expected benefit and the effect on model-recommended allocation. These logs can later be used to improve the algorithm and to defend allocation decisions during regulatory or institutional review.

Empirical Evaluation Strategy

A publishable empirical evaluation should not begin with grand claims about reducing national shortages. It should begin with a transparent comparison against operational baselines. The most defensible baselines are reorder-point policy, days-of-supply threshold policy, proportional allocation, first-come, first-served allocation, and manual pharmacist-led redistribution. The proposed integrated model should be evaluated against these baselines using retrospective data before being tested prospectively. Anything else is marketing, not science.

The retrospective study should reconstruct historical decision points. At each decision point, only information available at that time should be used. This prevents look-ahead bias, a common and rather embarrassing flaw in applied analytics papers. The model should then forecast demand, estimate stockout risk and recommend allocations. Outcomes should be compared with observed stockout exposure, emergency purchasing, substitution burden and allocation fairness. Where observed allocation records are unavailable, simulation can be used, but the assumptions must be stated plainly. The prospective pilot should be designed as a stepped implementation rather than an uncontrolled technology deployment. Facilities can be grouped by size, rurality, service-line mix and baseline data readiness. The platform can then be introduced in phases while continuing to monitor comparable facilities under existing decision rules. This design is not as clean as a randomised trial, but it is often more feasible in supply chain operations. The key is to define endpoints before implementation begins.

REFERENCES

1. American Society of Health-System Pharmacists (ASHP). Drug Shortages Statistics. Available at: https://www.ashp.org/drug-shortages/shortage-resources/drug-shortages-statistics. Accessed 15 June 2026.
2. American Society of Health-System Pharmacists (ASHP). National Drug Shortages: 2024 Drug Shortages Survey / Report. 2025. Available at: https://www.ashp.org/-/media/assets/drug-shortages/docs/2024/2024-Drug-Shortages-Survey.pdf. Accessed 15 June 2026.
3. U.S. Food and Drug Administration (FDA). Drug Shortages. Available at: https://www.fda.gov/drugs/drug-safety-and-availability/drug-shortages. Accessed 15 June 2026.
4. U.S. Food and Drug Administration (FDA). Drug Shortages: Report to Congress, Calendar Year 2024. Available at: https://www.fda.gov/media/189325/download. Accessed 15 June 2026.
5. U.S. Government Accountability Office (GAO). Drug Shortages: HHS Should Implement a Mechanism to Coordinate Its Activities. GAO-25-107110. 9 April 2025.
6. U.S. Department of Health and Human Services, Office of the Assistant Secretary for Planning and Evaluation. Policy Considerations to Prevent Drug Shortages and Mitigate Supply Chain Vulnerabilities in the United States. 2024.
7. U.S. Food and Drug Administration. Supply Chain News, Reports and Publications. Includes HHS 2025-2028 Draft Action Plan for Addressing Shortages of Medical Products and Critical Foods. Available at: https://www.fda.gov/emergency-preparedness-and-response/supply-chain/supply-chain-news-reports-and-publications. Accessed 15 June 2026.
8. McGeeney JD, et al. Analysis of Drug Shortages, 2018-2023. NCBI Bookshelf. 2025.
9. IQVIA Institute for Human Data Science. Drug Shortages in the U.S. 2023. 2023.
10. Office of the Assistant Secretary for Preparedness and Response (ASPR). National Strategy for a Resilient Public Health Supply Chain. U.S. Department of Health and Human Services. 2021.