Journal of Physical Sciences: Theories, Methods and Applications
Volume 1, Issue 1 (2026), pp. 1–26
ISSN: xxxx-xxxx (print), xxxx-xxxx (online)
https://jpstma.org
DOI: https://doi.org/10.68839/jpstma.2026.zn2nxmmz

Research article

Multi-Stage Stochastic Healthcare Resource Allocation under Displacement-Driven Demand Uncertainty: Evidence from Nigeria

Chinedu Bright Okechukwu[1,*], Bright Okore Osu[2], Ijioma Onwukwe[3]

[1,2,3]*Department of Mathematics, Abia State University, Uturu, Abia State, Nigeria*

Email addresses

[1] cbokechukwu@waec.org.ng, [2] Osu.bright@abiastateuniversity.edu.ng, [3] onwukweijioma@abiastateuniversity.edu.ng

ABSTRACT

Healthcare systems affected by displacement must allocate scarce resources before demand is fully known. This study formulates a three-stage stochastic linear programme for primary healthcare and maternal and newborn health across 177 Nigerian local government areas and 348 supply hubs. Twelve demand states combine four empirical spatial archetypes with severity variation. Controlled comparisons distinguish stochastic initial planning, sequential adaptation and proportional-coverage restrictions. The baseline value of stochastic planning is approximately 0.000338% of expected loss; primary healthcare and maternal and newborn health coverage are 37.74% and 70.96%, respectively. Highly correlated demand vectors, exhausted budgets and shortage-dominated loss accompany the small information gains. Held-out-archetype tests preserve small advantages for stochastic initial plans. Capacity and demand perturbations have larger effects at the levels examined, but capacity expansion at fixed budgets can shift service away from maternal and newborn health. A matched equal-coverage restriction improves minimum local coverage while reducing aggregate service. These conditional planning results distinguish information flexibility from resource scarcity and distributional policy choices.

ARTICLE HISTORY

Submitted
21 October 2026

Accepted
24 October 2026

Published online
26 October 2026

CITE THIS ARTICLE

C. B. Okechukwu, B. O. Osu, and I. Onwukwe. Multi-Stage Stochastic Healthcare Resource Allocation under Displacement-Driven Demand Uncertainty: Evidence from Nigeria. *J. Phys. Sci., Theor. Methods Appl.*, 1(1), 2026, 1–26.

Keywords: Stochastic programming, healthcare allocation, displacement, sequential adaptation, resource scarcity.

Mathematics Subject Classification: 90C15, 90C05, 90B50

1 Introduction

Internal displacement changes healthcare needs while resources remain constrained. Reviews and Nigerian audits document unmet needs and uneven programme

* **Corresponding Author:** cbokechukwu@waec.org.ng

© 2026 by the authors. Published by the Journal of Physical Sciences: Theories, Methods and Applications (JPSTMA). This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

provision [6, 9]. Geographic accessibility also varies, requiring feasible links between populations and facilities [1, 19].

This study evaluates a three-stage stochastic linear programme for primary healthcare (PHC) and maternal and newborn health (MNH) across 177 Nigerian local government areas (LGAs) and 348 supply hubs. Twelve demand states combine four displacement-derived spatial archetypes with severity variation. Common-input comparisons distinguish stochastic initial planning, sequential adaptation, geographic allocation restrictions and resource changes. The contribution is empirical evaluation, not a new optimisation algorithm. The objective is risk-neutral; conditional value-at-risk (CVaR) is diagnostic. PHC visit-equivalents and MNH delivery-equivalents remain separate service measures. The quarterly model evaluates deployment rather than treatment timeliness.

1.1 Literature review and research positioning

Mathematical programming connects healthcare objectives to feasible expenditure and delivery choices [24, 25]. Chalabi et al. [7] incorporate uncertainty and information value into programme allocation. Stochastic programming links early commitments to later recourse through non-anticipativity [5, 23]. Its value requires common-environment evaluation, rather than comparing objectives from different demand settings [4, 10].

Healthcare applications represent different uncertainty mechanisms. Yao et al. [28] address activity durations and care sequencing; Yin and Büyüktahtakın [29] model epidemic progression with equity. Fattahi et al. [11] include sharing and patient transfers, while Baas et al. [2] coordinate hospital capacity and patient assignment. Kapungu et al. [17] examine stochastic health-insurance financing. These studies motivate staged decisions but do not establish information value for displacement-driven service flows over a fixed PHC/MNH network.

Humanitarian pre-positioning separates initial placement from response [21], while evacuation models connect facilities, budgets and network feasibility [8]. Here, hubs and admissible links remain fixed. Displacement evidence motivates spatial representation but does not identify universal clinical calibration coefficients. Accessibility describes potential reach, not utilisation or effective throughput.

Scenario trees represent dependence and information structure [15]. The companion scenario study [31] preserves observed spatial vectors while scaling severity, leaving unobserved relocation outside the scenario set. Recent humanitarian models test decisions beyond training scenarios [3, 14]. Our held-out tests are narrower: they evaluate initial commitments with conditional recourse rather than a complete policy for arbitrary unseen states.

Risk and equity require explicit definitions. Rockafellar and Uryasev [22] establish CVaR for general distributions; Yin et al. [30] incorporate it into ventilator allocation. Merely reporting CVaR does not make an objective risk-averse. Equal amounts, equal coverage and need-based priorities also express different equity preferences [18]. A matched equal-coverage experiment therefore complements the separate proportional-planning policy. Table 1 positions the study within this literature.

Table 1: Selected studies and their relationship to the present allocation problem

Study	Setting and method	Relevant contribution	Distinction in the present study
Chalabi et al. [7]	Healthcare programme allocation; two-stage stochastic optimisation	Uncertainty and information value in budget allocation	LGA-level physical network and sequential resource deployment
Yin and Büyüktahtakın [29]	Ebola in West Africa; multi-stage allocation with equity	Links disease uncertainty, resource decisions and equity	Empirical displacement-demand states rather than epidemic dynamics

Continued on next page

Table 1 (continued)

Study	Setting and method	Relevant contribution	Distinction in the present study
Fattahi et al. [11]	Pandemic systems; multi-stage planning and rolling horizon	Combines resource sharing, transfers and capacity options	Existing PHC/MNH network with a narrower allocation action set
Yao et al. [28]	Patient management; two-stage stochastic scheduling	Coordinates multiple resources under uncertain durations	Geographic service allocation rather than care sequencing
Baas et al. [2]	Regional hospitals; stochastic lookahead	Links room capacity and patient assignment with lead times	Archetype and severity information within quarterly planning
Lawal and Anyiam [19]; Averik et al. [1]	Nigerian primary healthcare; geospatial access analysis	Measures potential geographic reach of facilities	Chooses feasible allocations subject to access and capacity
Ekezie et al. [9]	Nigerian displaced-population settings; empirical audit	Documents uneven health-resource provision	Evaluates allocation policies using calibrated spatial demand

2 Materials and methods

2.1 Data and planning system

Appendix A documents the demand-source values and limitations. The calibration follows the companion scenario study [31]; it supplies planning inputs, not independently validated clinical requirements.

2.1.1 Study geography and data integration

The common panel contains 177 LGAs in Benue, Kaduna, Kano, Katsina, Kogi, Nasarawa, Niger, Plateau, Sokoto and Zamfara States across IOM DTM rounds R15–R18, giving 708 LGA-round observations. The fixed geographic set excludes differences caused solely by survey coverage.

DTM displacement records, GRID3 population and facility inputs, and administrative boundaries are joined through state–LGA keys. Nnanatu et al. [20] document the population release. LGA totals and population-weighted centroids define demand; supply is grouped by LGA and facility level.

Population remains fixed. Displacement modifies service intensity rather than adding a separate resident headcount; demand concerns displaced people and host communities sharing services.

2.1.2 Quarterly demand and scenario construction

The quarterly PHC coefficient is exactly 45,000,000/237,527,782. Its approximate value is 0.18945152276966068. Reported Q2 2025 BHCPF-facility visits and national population have different coverage bases, making this a planning anchor rather than observed utilisation [12, 27]. MNH uses the 2024 crude birth rate of 32.493 per 1,000 annually and exact state facility-delivery proportions [26]. Both displacement premiums are assumed to be 0.50.

Let P_i denote fixed LGA population and I_{ik} displaced population in archetype k. Primary demand and quarterly births are

$$\mu_{i,PHC,k} = \frac{45,000,000}{237,527,782}(P_i + 0.50I_{ik}) \tag{1}$$

$$B_i = P_i\frac{32.493}{4,000} \tag{2}$$

$$\mu_{i,MNH,k} = \frac{32.493}{4,000} FBR_{state(i)}(P_i + 0.50I_{ik}) \tag{3}$$

Here $FBR_{state(i)}$ is the exact NDHS 2024 facility-delivery proportion. Division by 4,000 converts annual births per 1,000 into quarterly births per person. Non-comparable vulnerability and maternal subgroup fields, including unavailable R16 maternal fields, are excluded rather than zero-imputed.

Table 2: Demand and cost calibration status

Input	**Model value**	**Provenance and interpretation**
PHC quarterly utilisation	$\frac{45{,}000{,}000}{237{,}527{,}782}$	Exact ratio of reported BHCPF-facility quarterly visits to the national population anchor; different coverage bases.
Displacement modifiers	0.50 for PHC and MNH	Assumed displacement service-intensity premiums; no separate IDP headcount addition.
Vulnerability modifiers	Excluded from primary demand	General vulnerability and maternal subgroup fields are not sufficiently comparable across rounds; R16 maternal data are unavailable.
Annual birth rate	32.493 per 1,000	Nigeria, 2024; annual national crude birth rate [26].
Facility-birth proportion	State specific	Exact NDHS 2024 Table 9.8 live-birth facility-delivery proportions; listed in Appendix A.
Severity multipliers	0.90, 1.00, 1.15	Planning stresses applied to complete archetypes.
Severity weights	0.20, 0.60, 0.20	Design weights; not estimated event frequencies.
Procurement cost	PHC 1.00–1.25; MNH 3.00–3.75	Normalised operating-cost indices, not observed prices.
Shortage penalty	PHC 10; MNH 25	Normative model weights; MNH alternatives 12.5 and 50 tested in Section 3.11.

The 708 primary rows expand to 4,248 scenario–LGA–resource records, with 354 records per scenario.

Each round supplies one complete spatial archetype. Multipliers 0.90, 1.00 and 1.15 give LOW, BASE and HIGH scenarios while preserving spatial proportions.

Archetypes each receive weight 0.25. Conditional severity weights 0.20, 0.60 and 0.20 give terminal weights 0.05, 0.15 and 0.05. These are planning assumptions, not event frequencies. The mean multiplier is 1.01, not BASE.

2.1.3 Supply hubs and spatial accessibility

The 12,245 represented facilities form 348 hubs: 177 primary, 144 secondary and 27 tertiary. These are planning aggregates; effective capacity is calibrated separately.

The fixed network contains 885 PHC and 1,416 MNH links, repeated over three periods to form 6,903 allocation-cost records.

The calibrated network uses distance and effective-speed proxies for accessibility. The regional capacity stress retains this network; route outages and changing travel conditions are not simulated. Figure 1 summarises this structure.

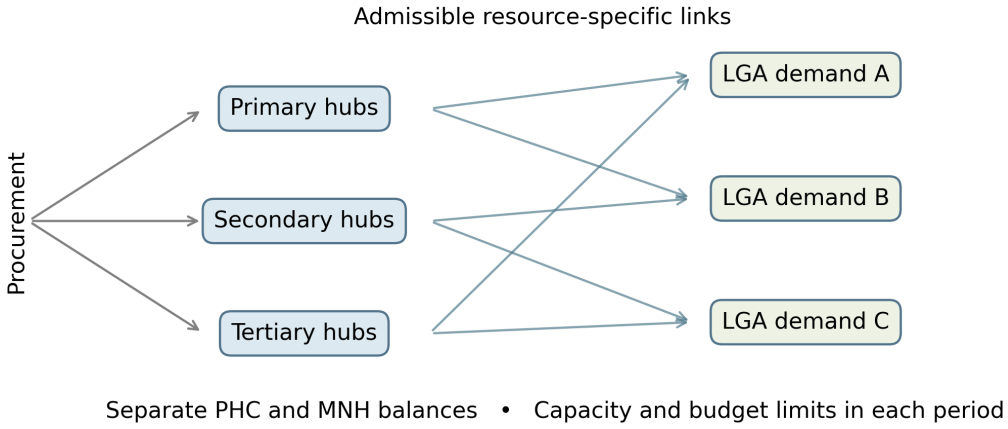


Figure 1: Conceptual allocation network. Procurement feeds aggregated hubs, and service flows only along admissible resource-specific links to LGA demand nodes. The illustrated links are schematic; the implemented network contains 885 PHC and 1,416 MNH links. Separate capacity, budget and cumulative inventory constraints apply

2.1.4 Capacity costs budgets and inventory

Capacity equals represented facility count times nominal throughput times readiness. Quarterly capacity is partitioned into period shares 0.40, 0.30 and 0.30, giving 2,088 hub–resource–period records. These are planning proxies, not observed throughput.

Normalised procurement costs are 1.00–1.25 per PHC unit and 3.00–3.75 per MNH unit. Allocation costs depend on link and period. Shortage penalties are 10 and 25, respectively, uniform across LGAs; no vulnerability-weighted objective is imposed.

Stage budgets are 90% of full-capacity procurement costs: 2,048,311.83, 1,689,857.26 and 1,920,292.34 cost-index units. Constraints cover procurement plus allocation, so 90% implies neither service coverage nor capacity utilisation. Values are not Naira estimates.

Initial inventory is zero. Procurement supports current or later service subject to cumulative stock, capacity and budget limits (Table 3).

Table 3: Planning system dimensions and baseline calibration

Component	**Specification**
Demand system	177 LGAs; PHC and MNH; twelve scenarios
Supply representation	348 hubs: 177 primary, 144 secondary and 27 tertiary
Admissible links	885 PHC arcs and 1,416 MNH arcs
Capacity stages	Shares 0.40, 0.30 and 0.30 of quarterly capacity
Initial stock	Zero in the baseline
Stage budgets	2,048,311.83; 1,689,857.26; 1,920,292.34 cost-index units
Shortage penalties	PHC: 10; MNH: 25 cost-index units per unmet unit

Continued on next page

Table 3 (continued)

Component	Specification
Procurement cost ranges	PHC: 1.00–1.25; MNH: 3.00–3.75 per unit
Empirical panel	Four DTM rounds; 708 LGA-round observations
Scenario records	4,248 records; $12 \times 177 \times 2$ (scenarios × LGAs × resources)
Represented facilities	12,245 facilities aggregated into 348 hubs

2.1.5 Information revelation and physical timing

The tree has one root, four archetype nodes and twelve terminal severity nodes. Decisions coincide across scenarios sharing an observed history.

All policies retain three physical periods and their resources. TS2 observes archetype but not severity; STATIC fixes each period across scenarios. Historical rounds are alternative archetypes, not successive periods (Figure 2).

Quarter-end demand balance imposes no interim service deadlines, waiting times or deterioration. Information stages therefore do not establish treatment timeliness.

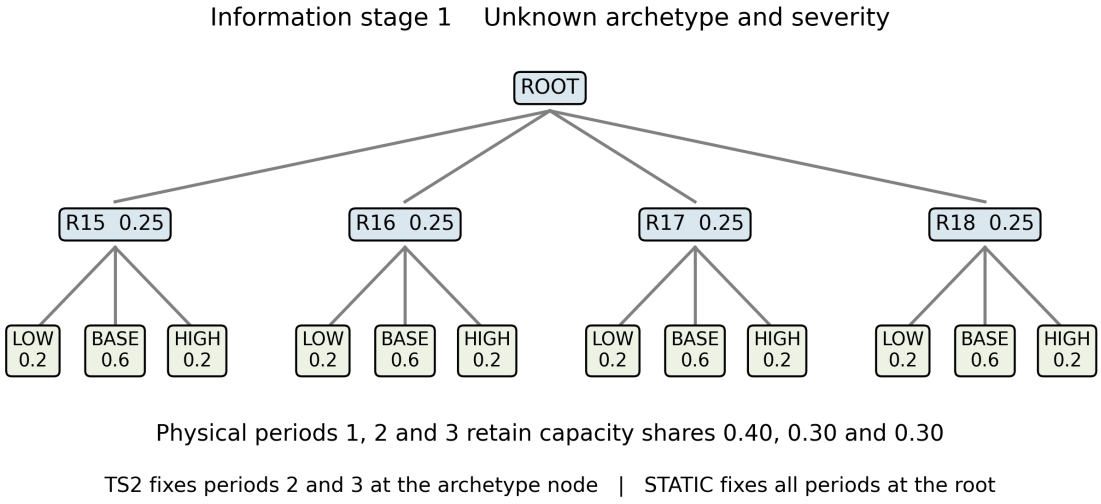


Figure 2: Scenario tree and physical timing. Four alternative archetypes each have weight 0.25; conditional LOW, BASE and HIGH weights are 0.20, 0.60 and 0.20. Stage 2 reveals the archetype and stage 3 reveals severity. Historical rounds are alternative states, not successive quarterly periods. Terminal weights are 0.05, 0.15 and 0.05 within each archetype

2.1.6 Input consistency and interpretation

Checks confirm unique keys, common scenario support, nonnegative demand, severity scaling, consistent weights, valid network endpoints and matching capacity/procurement keys. Capacity products and budget benchmarks reconcile. These establish internal consistency, not source validity.

2.2 Mathematical formulation

2.2.1 Sets parameters and decision variables

Let I be the set of 177 LGA demand nodes, J the 348 supply hubs, $R = \{PHC, MNH\}$ the resource classes and $T = \{1, 2, 3\}$ the physical periods. The admissible network is $A \subseteq I \times J \times R$. Define $J_{ir} = \{j : (i, j, r) \in A\}$ and $I_{jr} = \{i : (i, j, r) \in A\}$, so sums include only feasible links. A sum over an empty neighbour set is zero. The finite scenario set is Ω, with positive planning weights p_ω satisfying $\sum_{\omega\in\Omega} p_\omega = 1$.

The map $n_t(\omega)$ identifies the information node reached by scenario ω in period t. All scenarios share the root; scenarios belonging to the same archetype share a second-stage node; each terminal scenario has its own third-stage node. The map $k(\omega)$ identifies the archetype. Table 4 defines the remaining notation. Parameters other than demand are deterministic in the baseline; scenario-weight and resource perturbations are specified separately in Section 2.3.

Table 4: Parameters and variables of the allocation models

Symbol	Definition
D^{ω}_{ir}	Quarterly terminal demand in LGA i for resource r under scenario ω
K_{jrt}	Effective allocation capacity of hub j for resource r in period t
B_t	Expenditure limit for each information node in period t
I^0_{jr}	Initial inventory at hub j for resource r
$c^P_{jrt},\ c^A_{ijrt}$	Per-unit procurement and allocation cost coefficients
h_{ir}	Per-unit penalty for terminal unmet demand
x^{ω}_{jrt}	Quantity procured at hub j for resource r in period t
y^{ω}_{ijrt}	Quantity allocated from hub j to LGA i on admissible link (i, j, r) in period t
u^{ω}_{ir}	Terminal unmet demand in LGA i for resource r
$C_\omega,\ L_\omega$	Scenario operating cost and total loss, respectively
$\mathcal{F}_3,\ Z_3$	Feasible set and optimal expected loss of the three-stage benchmark

All variables are continuous and data finite and nonnegative. Superscript ω indexes scenarios, with $D^{\omega}_{ir} = D_{ir}(\omega)$. PHC/MNH balances stay separate; shared-node variables implement information restrictions.

2.2.2 Scenario costs and risk neutral objective

$$C_\omega = \sum_{t\in T}\left(\sum_{j\in J}\sum_{r\in R} c^P_{jrt}x^{\omega}_{jrt} + \sum_{(i,j,r)\in A} c^A_{ijrt}y^{\omega}_{ijrt}\right). \tag{4}$$

$$L_\omega = C_\omega + \sum_{i\in I}\sum_{r\in R} h_{ir}u^{\omega}_{ir}, \qquad \omega \in \Omega. \tag{5}$$

$$Z_3 = \min_{(x,y,u)\in\mathcal{F}_3} \sum_{\omega\in\Omega} p_\omega L_\omega. \tag{6}$$

Equations (7)–(12) define feasibility. Scenario weights count shared root costs once. No risk-aversion or vulnerability term enters the objective.

2.2.3 Demand capacity stock and budget constraints

Quarterly demand balance. For every $i \in I$, $r \in R$ and $\omega \in \Omega$,

$$\sum_{t \in T} \sum_{j \in J_{ir}} y^{\omega}_{ijrt} + u^{\omega}_{ir} = D^{\omega}_{ir}. \tag{7}$$

Equation (7) prevents over-service and resource substitution; it is a quarterly balance.

Effective capacity. For every $j \in J$, $r \in R$, $t \in T$ and $\omega \in \Omega$,

$$\sum_{i \in I_{jr}} y^{\omega}_{ijrt} \leq K_{jrt}. \tag{8}$$

Capacity limits allocation in the specified period. Unused service capacity cannot be carried into another period. Procurement alone does not create additional service capacity.

Cumulative stock conservation. For the same indices,

$$\sum_{\ell=1}^{t} \sum_{i \in I_{jr}} y^{\omega}_{ijr\ell} \leq I^{0}_{jr} + \sum_{\ell=1}^{t} x^{\omega}_{jr\ell}. \tag{9}$$

Stock carries forward but later procurement cannot fund earlier service. Residual stock may remain unused; transfers, storage limits, deterioration and salvage are absent.

Nodewise expenditure limits. For every $t \in T$ and $\omega \in \Omega$,

$$\sum_{j \in J} \sum_{r \in R} c^{P}_{jrt} x^{\omega}_{jrt} + \sum_{(i,j,r) \in A} c^{A}_{ijrt} y^{\omega}_{ijrt} \leq B_t. \tag{10}$$

Budgets apply once per information node, without probability scaling or carry-forward. Stock may carry forward under (9).

Variable domains. For all relevant indices,

$$x^{\omega}_{jrt} \geq 0, \qquad y^{\omega}_{ijrt} \geq 0 \ \ ((i,j,r) \in A), \qquad u^{\omega}_{ir} \geq 0. \tag{11}$$

No allocation variable is created for a link outside A. A disconnected LGA–resource pair therefore has all its demand recorded as unmet rather than making the baseline planning problem infeasible.

2.2.4 Nonanticipativity and feasibility

Decisions can use only the information available when they are made. For every pair $\omega, \omega' \in \Omega$ and period t,

$$n_t(\omega) = n_t(\omega') \ \Rightarrow \ x^{\omega}_{jrt} = x^{\omega'}_{jrt}, \quad y^{\omega}_{ijrt} = y^{\omega'}_{ijrt}. \tag{12}$$

These equalities cover all hubs and admissible links. Terminal shortages remain scenario-specific.

Lemma 1 Feasibility and demand-compatible early allocations. Under the stated nonnegative-data assumptions, the benchmark has a feasible solution and a finite attained optimum. Nevertheless, an arbitrary fixed early allocation need not admit feasible recourse. In particular, any feasible policy satisfies, at a node n in period t,

$$Q^{t}_{ir}(n) := \sum_{\ell=1}^{t} \sum_{j \in J_{ir}} y^{\omega}_{ijr\ell} \leq \min_{\omega' \in \Omega(n)} D^{\omega'}_{ir}, \tag{13}$$

where $\Omega(n) = \{\omega' : n_t(\omega') = n\}$ and $\omega \in \Omega(n)$. The accumulated allocation is common across these descendants because their histories through period t coincide.

Proof. Zero procurement/allocation with unmet demand equal to demand is feasible. Its finite loss and the nonnegative objective imply an attained finite LP optimum. Nonnegative later allocation and shortages in (7) bound early service by every descendant demand, proving (13). Conversely, zero later decisions complete any physically feasible history satisfying the bound. Shortages cannot absorb excess early allocation.

2.2.5 Restricted adaptation and policy ordering

The two-information-stage policy, denoted $TS2$, retains all three physical periods but suppresses the final severity update. In addition to (7)–(12), impose the following equalities for $t \in \{2, 3\}$:

$$k(\omega) = k(\omega') \;\Rightarrow\; x^{\omega}_{jrt} = x^{\omega'}_{jrt}, \quad y^{\omega}_{ijrt} = y^{\omega'}_{ijrt}. \tag{14}$$

Thus third-period quantities can differ across archetypes but cannot differ among LOW, BASE and HIGH descendants of the same archetype. The static staged policy, denoted $STATIC$, instead imposes, for all periods and all scenario pairs,

$$x^{\omega}_{jrt} = x^{\omega'}_{jrt}, \qquad y^{\omega}_{ijrt} = y^{\omega'}_{ijrt}. \tag{15}$$

$STATIC$ permits different planned quantities in different physical periods; it removes their dependence on the scenario. Both comparators retain scenario-specific shortages and minimise the same expected-loss objective. Their feasible sets and optimum values are denoted $\mathcal{F}_2, Z_2$ and $\mathcal{F}_0, Z_0$, respectively.

Proposition 1 Ordering under restricted information. If all three policies share demand, probabilities, cost coefficients, capacity, inventory, budgets and accessibility, then

$$\mathcal{F}_0 \subseteq \mathcal{F}_2 \subseteq \mathcal{F}_3, \qquad Z_3 \leq Z_2 \leq Z_0. \tag{16}$$

Proof. STATIC satisfies TS2 restrictions, which imply the benchmark restrictions. Physical constraints and losses remain common. Minimisation over nested feasible sets gives (16); Lemma 1 ensures finite optima.

The inequalities guarantee neither strictly positive benefits nor improvements in each service. Relative gaps divide by $Z_3 > 0$.

2.2.6 Expected value model and stochastic first stage value

The deterministic expected-value model replaces terminal demand by its probability-weighted mean:

$$\overline{D}_{ir} = \sum_{\omega \in \Omega} p_{\omega} D^{\omega}_{ir}. \tag{17}$$

Let $\mathcal{F}_{EV}$ denote (7)–(11) with D^{ω}_{ir} replaced by $\overline{D}_{ir}$ and all scenario indices removed. It retains the three periods, network and physical resource limits. Writing $C(x, y)$ for equation (4) without scenario indices, solve

$$Z_{EV} = \min_{(x,y,\overline{u}) \in \mathcal{F}_{EV}} [C(x, y) + \sum_{i \in I} \sum_{r \in R} h_{ir} \overline{u}_{ir}]. \tag{18}$$

Choose an optimal expected-value plan and denote its first-period decisions by $\widehat{x}_{jr1}$ and $\widehat{y}_{ijr1}$. Evaluate those decisions by restoring the complete stochastic tree and solving

$$EEV = \min\{\sum_{\omega \in \Omega} p_{\omega} L_{\omega} : (x, y, u) \in \mathcal{F}_3,\; x^{\omega}_{jr1} = \widehat{x}_{jr1},\; y^{\omega}_{ijr1} = \widehat{y}_{ijr1}\}. \tag{19}$$

Only first-period procurement and allocation are fixed; later recourse remains adaptive.

Given a first-stage plan satisfying the physical constraints, feasible continuation exists in this model if and only if

$$\sum_{j \in J_{ir}} \widehat{y}_{ijr1} \leq \min_{\omega \in \Omega} D^{\omega}_{ir} \qquad \text{for every } i, r. \tag{20}$$

Lemma 1 gives necessity; zero later decisions prove sufficiency. Failure implies infeasible EEV, not finite VSS.

With $RP = Z_3$, define

$$VSS = EEV - Z_3, \qquad VSS_{rel} = 100\frac{EEV - Z_3}{Z_3} \quad (Z_3 > 0)\,. \tag{21}$$

Feasible EEV restricts the benchmark, so VSS is nonnegative. The deterministic objective Z_{EV} is not EEV. VSS depends on the selected mean-demand optimum; alternatives are not exhaustively enumerated.

2.2.7 Uniform proportion planning and scenario evaluation

The proportional comparator uses scenario-independent procurement and allocation. Let a_{ir} be total planned service and ρ_r the common planned fraction of mean demand for resource r. Impose

$$a_{ir} := \sum_{t\in T}\sum_{j\in J_{ir}} y_{ijrt} = \rho_r \overline{D}_{ir}, \qquad 0 \le \rho_r \le 1. \tag{22}$$

The model retains (8)–(10) without scenario indices and requires nonnegative procurement and allocation. There is no shortage variable in its planning problem. Equation (22) replaces scenario demand balance at the planning stage. The implemented objective is

$$\min_{x,y,\rho}[C(x,y) - M\sum_{r\in R}\rho_r], \qquad M = 10^8. \tag{23}$$

The finite weight rewards both coverage fractions equally but does not establish exact two-pass lexicographic optimality.

For the selected plan, evaluated service s^{ω}_{ir}, unmet demand and unused planned service w^{ω}_{ir} are

$$\begin{aligned} s^{\omega}_{ir} &= \min\{a_{ir}, D^{\omega}_{ir}\},\\ u^{\omega}_{ir} &= \max\{D^{\omega}_{ir} - a_{ir}, 0\},\\ w^{\omega}_{ir} &= \max\{a_{ir} - D^{\omega}_{ir}, 0\}. \end{aligned} \tag{24}$$

The reported expected loss charges the full planned operating cost and scenario-dependent shortage penalties:

$$Z_{PROP} = C(x,y) + \sum_{\omega\in\Omega} p_\omega \sum_{i\in I}\sum_{r\in R} h_{ir} u^{\omega}_{ir}. \tag{25}$$

PROP charges unused planned service without refund or reassignment. It changes the objective and service accounting, so its gap is not a pure information effect.

2.2.8 Resource expansion and the limits of monotonicity

Proposition 2 Monotonicity under resource expansion. Hold demand, weights, costs, accessibility and information restrictions fixed. For the benchmark or either restricted-information model, let $Z\left(K, B, I^0\right)$ denote optimal expected loss. If resource bounds increase componentwise, then

$$K' \ge K,\ B' \ge B,\ (I^0)' \ge I^0 \quad \Rightarrow \quad Z\left(K', B', (I^0)'\right) \le Z\left(K, B, I^0\right). \tag{26}$$

Proof. Every originally feasible solution remains feasible: only the right-hand sides of capacity, budget and stock inequalities are relaxed. Demand balance, nonnegativity, information restrictions and the objective are unchanged. The new feasible set contains the old one, so its minimum cannot be larger.

Only optimal expected loss is ordered. Individual service outcomes, changed probabilities/demand and PROP are outside this guarantee.

2.2.9 Deterministic equivalent and implementation correspondence

Let N_t be the information nodes in period t, and define $q_n = \sum_{\omega:n_t(\omega)=n} p_\omega$ for $n \in N_t$. Write $g_t\,(x(n), y(n))$ for the period-t procurement and allocation cost inside the brackets of (4). The equivalent node-indexed objective is

$$\min[\sum_{t\in T}\sum_{n\in N_t} q_n g_t\,(x(n), y(n)) + \sum_{\omega\in\Omega} p_\omega \sum_{i\in I}\sum_{r\in R} h_{ir} u_{ir}^{\omega}]. \tag{27}$$

Collecting shared path costs gives node weights: one at the root, archetype probabilities at intermediate nodes and terminal probabilities at leaves. Physical constraints occur once per node; cumulative stock follows ancestor paths.

2.3 Evaluation and computational design

2.3.1 Policy comparisons and common accounting

Let v denote an evaluated policy and L_ω^v its realised base loss in scenario ω, calculated using the procurement, allocation and shortage accounting of Section 2.2. Its expected loss is

$$Z^v = \sum_{\omega\in\Omega} p_\omega L_\omega^v. \tag{28}$$

The $E00$ value is $Z^{E00} = Z_3$. Absolute and relative excess losses are reported as

$$\Delta Z^v = Z^v - Z^{E00}, \qquad G^v = 100\frac{Z^v - Z^{E00}}{Z^{E00}}. \tag{29}$$

PROP retains baseline evaluation weights but clips planned service at demand, charges its full costs and does not reassign unused service. Its gap combines objective and allocation-rule effects.

2.3.2 Resource coverage and the composite planning index

For each resource r, define expected demand $\overline{D}_r = \sum_\omega p_\omega \sum_i D_{ir}(\omega)$ and expected unmet demand $\overline{U}_r^v = \sum_\omega p_\omega \sum_i u_{ir}^v(\omega)$. When $\overline{D}_r > 0$, resource-specific coverage is

$$Cov_r^v = \frac{\overline{D}_r - \overline{U}_r^v}{\overline{D}_r} = 1 - \frac{\sum_\omega p_\omega \sum_i u_{ir}^v(\omega)}{\sum_\omega p_\omega \sum_i D_{ir}(\omega)}. \tag{30}$$

PROP shortages follow clipped service. Zero-demand aggregates would have undefined coverage.

An auxiliary pooled planning index is defined by

$$Cov_{comp}^v = 1 - \frac{\sum_\omega p_\omega \sum_{i,r} u_{ir}^v(\omega)}{\sum_\omega p_\omega \sum_{i,r} D_{ir}(\omega)}. \tag{31}$$

This pooled index combines unlike service units, depends on measurement scale and is not patient coverage. Separate PHC/MNH ratios are primary.

Both measures are ratios of expectations. To make the distinction explicit, let $d_{r\omega} = \sum_i D_{ir}(\omega) > 0$ and $c_{r\omega}^v = 1 - \sum_i u_{ir}^v(\omega)/d_{r\omega}$. Then

$$Cov_r^v = \sum_\omega \widetilde{p}_{r\omega} c_{r\omega}^v, \qquad \widetilde{p}_{r\omega} = \frac{p_\omega d_{r\omega}}{\sum_{\omega'} p_{\omega'} d_{r\omega'}}. \tag{32}$$

Coverage is demand-weighted, not generally the probability-weighted mean of scenario ratios. Coverage differences use percentage points.

2.3.3 Tail loss under a finite planning distribution

$CVaR$ is calculated from the realised base losses of each evaluated policy, using the same scenario weights as its reported expected loss. Following Rockafellar and Uryasev [22], for $0 < \alpha < 1$,

$$CVaR_\alpha\left(L^v\right) = \min_{\eta\in\mathbb{R}}\{\eta + \frac{1}{1-\alpha}\sum_{\omega} p_\omega[L^v_\omega - \eta]_+\}, \tag{33}$$

Here $[a]_+ = \max(a, 0)$. Threshold minimisation evaluates a fixed policy at $\alpha = 0.95$; it does not change allocation.

Sort losses $L^v_{(1)} \geq \cdots \geq L^v_{(N)}$ and their probabilities. For $N = 12$, successive upper-tail masses are

$$a_k = \min\{p_{(k)}, [1-\alpha-\sum_{j<k} a_j]\ \}_+, \qquad CVaR_\alpha\left(L^v\right) = \frac{\sum_{k=1}^{N} a_k L^v_{(k)}}{1-\alpha}. \tag{34}$$

Boundary scenarios contribute only the required fraction, giving total mass $1 - \alpha$. At baseline each probability is at least 0.05, so CVaR equals maximum loss. Under reweighting, this holds only when maximal-loss mass remains at least 0.05. Loss ranking, not HIGH labels, defines the tail; weights are not estimated rare-event frequencies.

2.3.4 Parameter sensitivity and controlled changes

Table 5 specifies the one-factor changes. Write a sensitivity case as e and let w^e_k and v^e_s be its archetype and conditional severity weights. For scenario $\omega = (k, s)$, capacity multiplier κ_e and demand multiplier δ_e, the inputs are

$$p^e_{ks} = \frac{w^e_k v^e_s}{\sum_{k',s'} w^e_{k'} v^e_{s'}}, \qquad K^e_{jrt} = \kappa_e K_{jrt}, \qquad D^e_{ir}(\omega) = \delta_e D_{ir}(\omega). \tag{35}$$

Probability cases change weights only; capacity cases scale all hub capacities; demand cases scale all terminal demands, preserving spatial proportions.

Table 5: Re-optimisation cases and controlled inputs

Case	**Changed input**	**Inputs retained at baseline**
$\mathrm{ROB}_{\mathrm{BASE}}$	No change	All baseline inputs
P_{HIGH}	LOW, BASE, HIGH weights 0.10, 0.50, 0.40	Archetype weights and physical system
P_{LOW}	LOW, BASE, HIGH weights 0.35, 0.55, 0.10	Archetype weights and physical system
P_{RECENT}	R15, R16, R17, R18 weights 0.10, 0.20, 0.30, 0.40	Severity weights and physical system
K80 and K120	Capacity multipliers 0.80 and 1.20	Demand, weights, budgets, costs and initial inventory
D90 and D110	Demand multipliers 0.90 and 1.10	Capacity, weights, budgets, costs and initial inventory

Links, costs and penalties remain fixed. Capacity/demand changes do not reset budgets; every case is re-optimised under its own inputs.

2.3.5 Effect sizes and limits of sensitivity interpretation

For case e, let Z_e and $Cov_{comp,e}$ be the re-optimised outcomes. Relative to ROB_{BASE}, denoted by subscript 0, the reported changes are

$$\Delta_e^{loss} = 100\frac{Z_e - Z_0}{Z_0}, \qquad \Delta_e^{cov} = 100\,(Cov_{comp,e} - Cov_{comp,0})\,. \tag{36}$$

Loss changes are percentages; coverage changes are percentage points. ROB_BASE reproduces E00.

Reweighting changes valuation and decisions; demand scaling also changes denominators. These experiments do not test a committed policy. Capacity expansion guarantees only weak expected-loss improvement.

2.3.6 Computational implementation and reproducibility

Python sparse deterministic-equivalent programmes use SciPy and HiGHS [16], shared-node variables and reconstructed scenario costs.

The 42 solves include thirteen common allocation/sensitivity runs, 25 baseline, regional, penalty and held-out runs, and four PROP, duplicated-mean and EQC checks. Repeated baselines are consistency checks, not independent evidence.

2.3.7 Numerical verification and reporting boundaries

Computations use Python 3.12.14, NumPy 2.3.5, SciPy 1.17.0 and HiGHS 1.8.0 on Linux x86_64. Primal/dual feasibility tolerances are 10^{-8}; other options retain defaults. Logs record status, runtime, residuals and objective reconstruction.

All solves terminate optimally. The baseline has 55,197 variables, 4,248 equalities, 23,681 inequalities and 312,015 nonzeros; runtime is 20.07 seconds. Maximum equality, inequality and stationarity residuals are 5.82×10^{-11}, 1.40×10^{-9} and 1.07×10^{-14}; cost reconstruction differs by 1.49×10^{-7}.

Across the common allocation/sensitivity block, equality/inequality errors are at most $1.16 \times 10^{-10}/3.49 \times 10^{-9}$; reconstruction differs by at most 4.75×10^{-6}. Additional checks are feasible below 10^{-7}, with objective discrepancies below 10^{-5}, substantially smaller than policy gaps. These checks validate numerical consistency, not calibration or statistical significance.

2.3.8 Validation and spatial outcome design

Scenario dependence uses Pearson correlation, cosine similarity and Euclidean distance between 177-LGA vectors by resource. Cross-archetype summaries exclude within-round pairs. Spatial coefficients of variation use population standard deviation divided by mean.

EQC retains the E00 objective, physical system and information structure, adding equal cumulative service-to-demand ratios across positive-demand LGAs in each resource/scenario. Fixed demands make ratio equality to a reference LGA linear. Unlike PROP, EQC neither changes the objective nor clips unused service.

Each held-out test trains stochastic and mean-demand initial plans on three archetypes with renormalised weights. First-period decisions are fixed on the excluded archetype; later recourse uses severity weights 0.20, 0.60 and 0.20. A test-aware plan can also change period 1, providing a lower bound. Shared calibration and nonchronological folds imply conditional transfer, not forecasting.

Regional stresses increase included Kaduna demand by 25% or halve its hub capacity, holding other inputs fixed. MNH penalties change from 25 to 12.5 or 50 while PHC remains 10. Different penalty objectives have different valuation scales; resource coverage is the main comparison.

Mapped LGA coverage is expected service divided by expected demand. Mean, median, minimum, tenth percentile, IQR and Gini give LGAs equal weight. Gini is pairwise absolute differences divided by twice the LGA count times total coverage. Maps show one optimum; alternatives may change local allocations.

3 Results and discussion

3.1 Baseline performance and comparison scale

The risk-neutral benchmark has expected loss of 81.043572 million cost units. PHC coverage is 37.743% and MNH coverage is 70.959%; these distinct quantities are the principal service outcomes. Baseline 95% CVaR is 98.192488 million, approximately 21.160% above expected loss. The auxiliary composite index is 38.200306% and its minimum across scenarios is 33.536950%. Section 3.8 reports the cost and utilisation diagnostics underlying this allocation.

Table 6 retains precision to reveal small model gaps, not to imply equivalent calibration accuracy. EQC appears separately in Section 3.9.

Table 6: Baseline policy performance

Policy	**Expected loss (million)**	**Composite coverage (%)**	$\mathrm{CVaR}_{0.95}$ **(million)**	**Excess loss (%)**
E00	81.043572	38.200306	98.192488	0.000000
EEV	81.043845	38.200215	98.193500	0.000338
TS2	81.063988	38.193546	98.231877	0.025192
STATIC	81.065253	38.193127	98.232687	0.026753
PROP	108.641483	12.537687	125.808917	34.053177

3.2 Stochastic first stage planning and sequential adaptation

Mean-demand optimisation gives 81,041,218.18 cost units. Its fixed-first-stage evaluation gives EEV 81,043,845.30 versus E00 81,043,571.63: VSS is 273.678 units (0.000338%). Subtracting the mean-demand objective instead compares different environments.

EEV lowers composite coverage by 0.000091 points and raises CVaR by 1,012.794 units. Implementation costs and calibration uncertainty govern practical significance.

Removing final-severity adaptation through TS2 increases expected loss by 20,416.420 cost units, or 0.025192%. Making all three physical periods scenario-independent through STATIC increases expected loss by 21,681.235, or 0.026753%. Composite coverage falls by 0.006760 and 0.007179 percentage points, respectively. Figure 3 plots differences from E00 so that these small gaps remain visible without truncating bars of absolute loss.

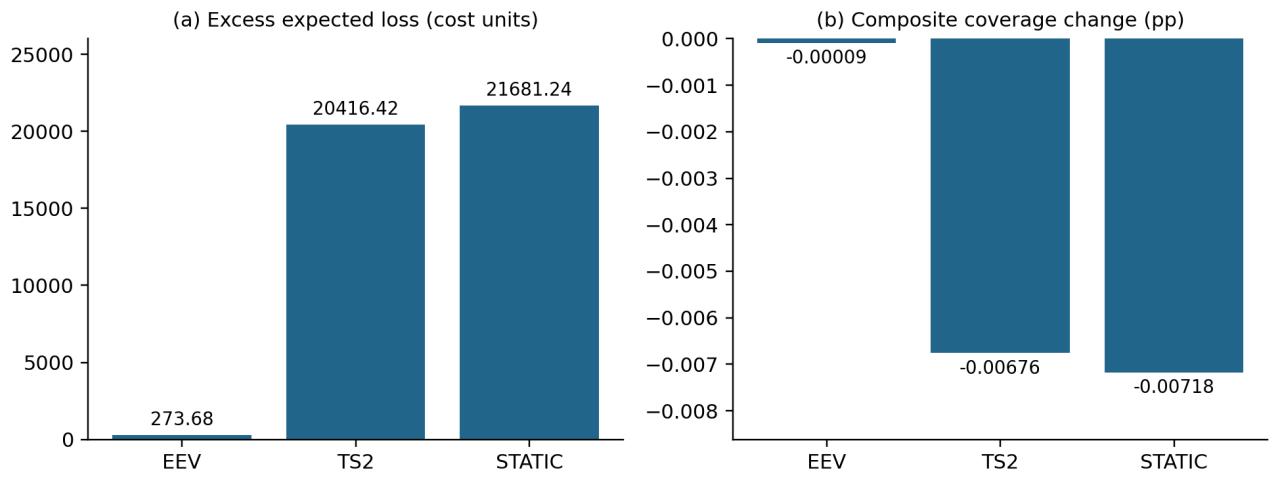


Figure 3: Information-related policy comparisons relative to $E00$. Panel (a) shows excess expected loss; Panel (b) shows the change in composite planning coverage. Each panel has its own units. EEV fixes mean-demand first-stage decisions, $TS2$ removes final-severity adaptation, and $STATIC$ removes scenario dependence from all physical decisions. Values report the evaluated policy outcomes

Archetype recourse improves STATIC loss by 1,264.815 units; severity recourse improves TS2 by 20,416.420. Together they equal STATIC–E00; VSS is separate and cannot be added.

A duplicated-branch mean-demand representation reproduces EEV of 81,043,845.304, but does not exhaust alternative mean-demand optima.

3.3 Performance of the uniform proportion rule

PROP plans PHC/MNH fractions of 11.983206% and 52.237344%. Realised aggregate coverage matches at reported precision; pooled coverage is 12.537687%, loss 108.641483 million and CVaR 125.808917 million.

PROP exceeds E00 loss by 34.053177% and reduces pooled coverage by 25.662619 points (Figure 4).

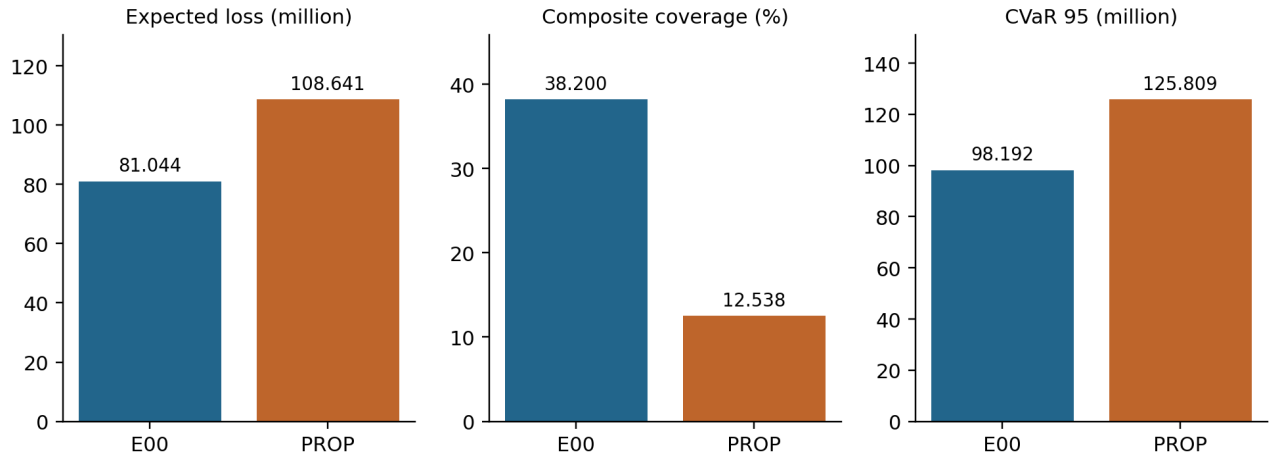


Figure 4: $E00$ and the implemented proportional policy. The panels report expected loss, composite planning coverage and $CVaR_{0.95}$. Lower loss is preferable; higher composite coverage indicates a larger pooled planning fraction. The comparison includes differences in planning objective and unused-service accounting as well as the uniform-proportion restriction; it is not a pure information-value comparison

The PROP gap combines proportionality, its planning objective and charged unused service. It neither isolates local bottlenecks nor establishes ethical superiority.

3.4 Sensitivity to scenario weights

Table 7 and Figure 5 report re-optimised sensitivity against ROB_BASE, which reproduces E00.

Table 7: Re-optimised expected loss and composite coverage

Case	Expected loss (million)	Composite coverage (%)	Loss change (%)	Coverage change (pp)	PHC coverage (%)	MNH coverage (%)
$\mathrm{ROB}_{\mathrm{BASE}}$	81.043572	38.200306	+0.000	+0.000	37.743	70.959
P_{HIGH}	85.921209	36.746895	+6.019	-1.453	36.305	68.390
P_{LOW}	77.386018	39.367995	-4.513	+1.168	38.898	73.008
P_{RECENT}	81.062948	38.194272	+0.024	-0.006	37.737	70.950
K80	88.832132	30.763294	+9.610	-7.437	30.194	71.509
K120	80.367262	40.226266	-0.835	+2.026	40.788	0.000
D90	68.737255	42.435845	-15.185	+4.236	41.936	78.195
D110	93.360400	34.731700	+15.198	-3.469	34.312	64.810

HIGH reweighting raises loss by 6.019% and lowers pooled coverage by 1.453 points; LOW lowers loss by 4.513% and raises coverage by 1.168 points. RECENT raises loss by 0.024%.

Probability effects combine reweighting and re-optimisation; RECENT is not a forecast or estimated trend.

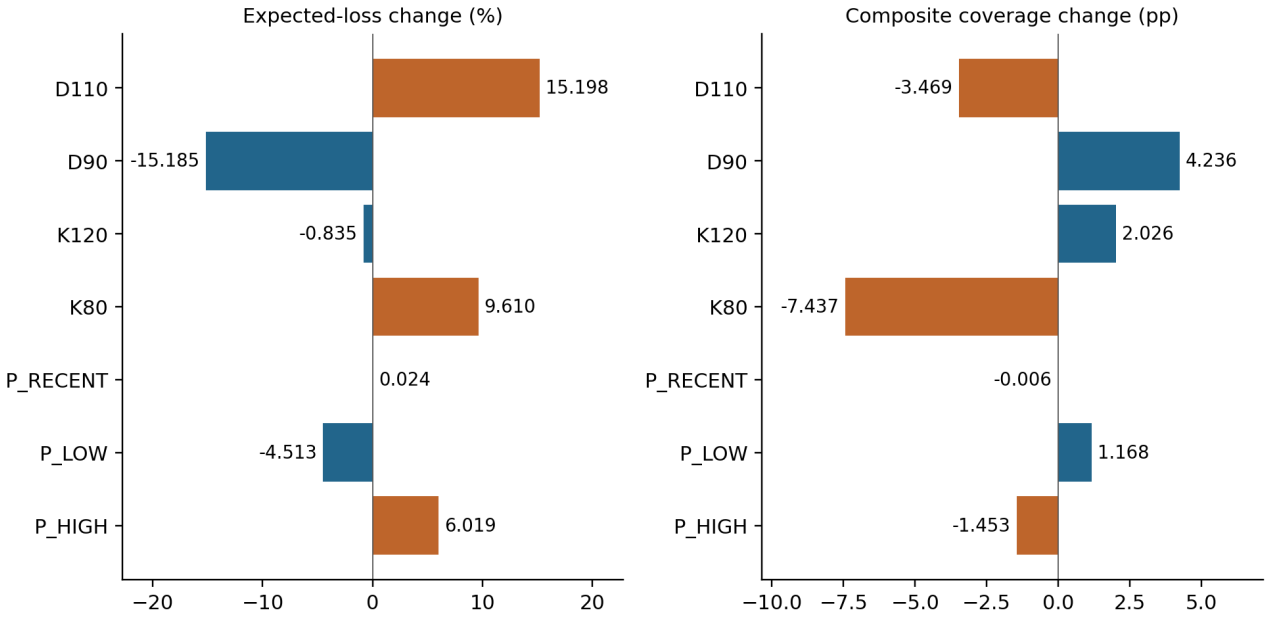


Figure 5: Re-optimised sensitivity relative to ROB_{BASE}. Expected-loss changes are percentages and composite-coverage changes are percentage points. Orange marks an adverse change in the displayed aggregate metric and blue a favourable change. The perturbations have different magnitudes and meanings, so the bars are not standardised elasticities or a causal ranking of input importance

3.5 Capacity and demand sensitivity

K80 raises loss by 9.610% and lowers pooled coverage by 7.437 points. K120 lowers loss by 0.835% and raises coverage by 2.026 points, consistent with Proposition 2.

Capacity effects are asymmetric under fixed budgets. K120 raises PHC coverage to 40.788% but eliminates MNH service. Expected-loss improvement need not benefit both services; mandatory floors require a separate comparison.

D90 lowers loss by 15.185%; D110 raises it by 15.198%. PHC/MNH coverage is 41.936%/78.195% and 34.312%/64.810%, respectively. Ratios reflect changed denominators; uniform scaling does not test relocation.

3.6 Tail loss and minimum scenario coverage

Table 8 reports CVaR and minimum pooled scenario coverage. The latter is neither minimum LGA coverage nor necessarily associated with maximum loss.

Table 8: Tail loss and minimum scenario coverage under sensitivity

Case	$\mathrm{CVaR}_{0.95}$ **(million)**	**CVaR change (%)**	**Minimum scenario coverage (%)**
$\mathrm{ROB}_{\mathrm{BASE}}$	98.192488	+0.000000	33.536950
P_{HIGH}	98.192404	-0.000085	33.536975
P_{LOW}	98.178937	-0.013800	33.536901
P_{RECENT}	98.192488	+0.000000	33.536950
K80	105.900839	+7.850245	27.037325
K120	97.531951	-0.672696	35.311917
D90	84.158969	-14.291846	37.258809
D110	112.233104	+14.299074	30.489918

HIGH and LOW change CVaR by only -0.000085% and -0.013800%, despite larger mean-loss changes. Minimum scenario coverage remains approximately 33.537% (Figure 6).

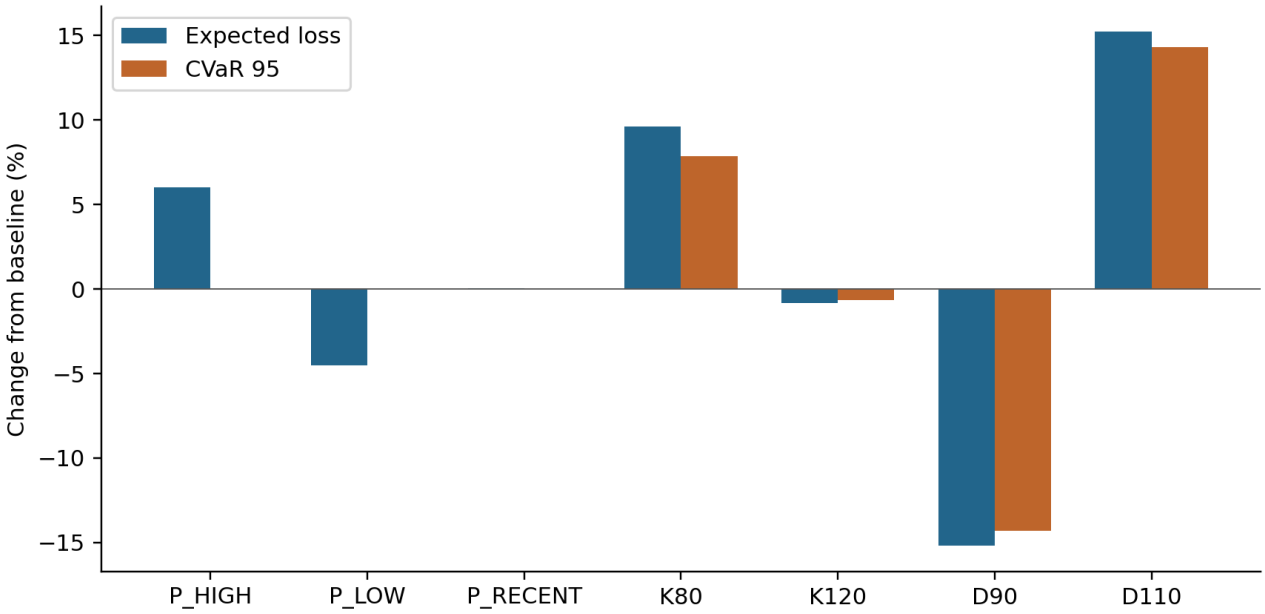


Figure 6: Expected-loss and tail-loss responses to re-optimised parameter changes. Each percentage uses its own baseline metric as denominator. Table 8 reports near-zero tail changes more precisely. CVaR is a post-solution diagnostic, and reweighted cases use the fractional-tail calculation in Section 2.3.

Capacity/demand changes have larger tail effects (Table 8). These diagnostics do not establish risk aversion or performance outside the scenario support.

3.7 Interpretation of the comparative evidence

Information restrictions, allocation rules and resource changes answer different questions. Their gaps cannot be combined into one measure of stochastic value.

3.8 Why the value of information is small

Shortages contribute 93.018% of baseline loss, procurement 6.614% and allocation 0.368%. All node budgets are exhausted; expected PHC/MNH capacity utilisation is 86.965%/68.278%. Only 17.645% of allocation variables exceed 10^{-7} units.

The 17,314 inequalities binding within 10^{-6} include zero-capacity and redundant rows, not independent causal bottlenecks. Budgets can bind while capacity remains unused.

Table 9: Cross-archetype similarity of demand vectors

Resource	Pearson correlation	Cosine similarity	Spatial CV
PHC	0.99992508–0.99998979	0.99997803–0.99999686	0.643967–0.644976
MNH	0.99994693–0.99999419	0.99997444–0.99999714	0.923883–0.926670

Cross-archetype correlations exceed 0.99992508 for PHC and 0.99994693 for MNH (Table 9; Figure 7). Positive scaling gives unit within-archetype correlation/cosine similarity. The common population base partly explains alignment, without establishing homogeneous displacement.

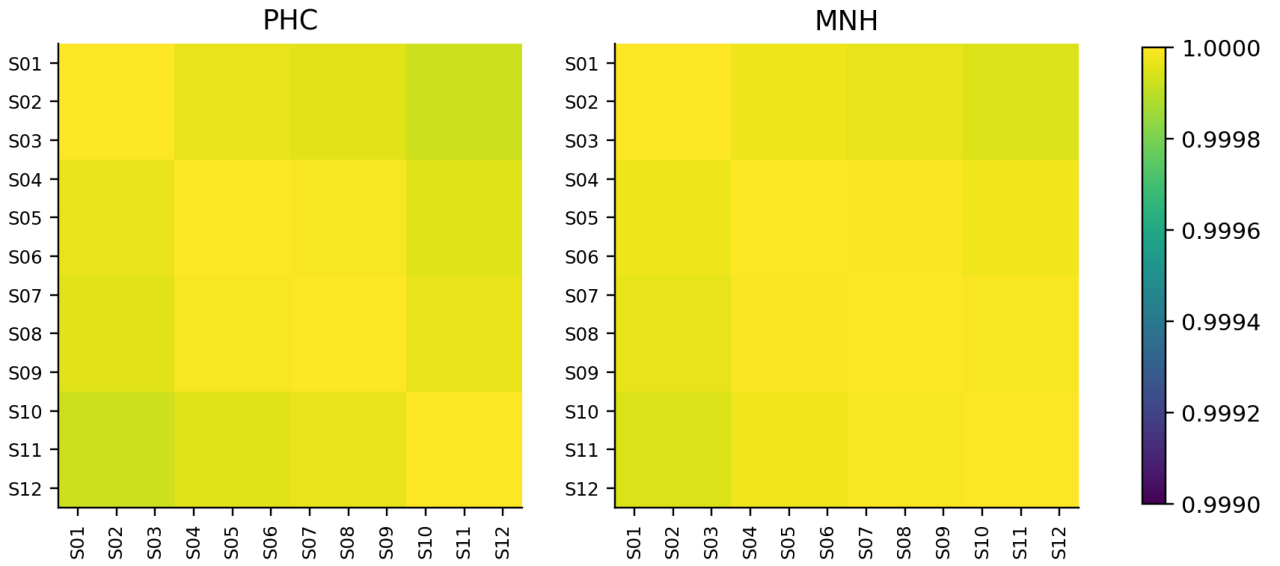


Figure 7: Pearson correlations across the twelve scenario demand vectors, separately for PHC (a) and MNH (b). S01–S03 correspond to R15, S04–S06 to R16, S07–S09 to R17 and S10–S12 to R18, each ordered LOW, BASE and HIGH. The colour scale runs from 0.999 to 1.000 to reveal the small differences. Positive severity scaling leaves within-archetype correlations unchanged

3.9 Resource coverage and the matched proportional restriction

Table 10: Matched comparison under a common expected loss objective

Policy	Loss (million)	PHC coverage (%)	MNH coverage (%)	Excess loss (%)
E00	81.043572	37.742762	70.959458	0.000000
EQC	108.641755	11.983206	52.237345	34.053513

EQC raises loss by 34.054% to 108.641755 million, with PHC/MNH coverage 11.983206% and 52.237345%. Only equal coverage is added, isolating its conditional cost. PROP differs by 272.006 units despite similar aggregate coverage. EQC node budgets are only 0.400–86.068% utilised.

Table 11: Distribution of LGA expected coverage ratios

Policy	Class	Mean (%)	Median (%)	Min (%)	P10 (%)	IQR (pp)	Gini
E00	PHC	46.977	43.080	2.407	16.723	35.180	0.29800931
EQC	PHC	11.983206	11.983206	11.983198	11.983204	0.000000	0.00000007
E00	MNH	82.292	98.247	6.086	44.689	30.277	0.14994316
EQC	MNH	52.237344	52.237344	52.237307	52.237340	0.000002	0.00000006

EQC raises minimum PHC/MNH coverage from 2.407% and 6.086% to approximately 11.983% and 52.237%, reducing aggregate service and dispersion (Table 11). Expected ratios can differ slightly because scenario demand weights vary by LGA. Figures 8–9 map one E00 optimum.

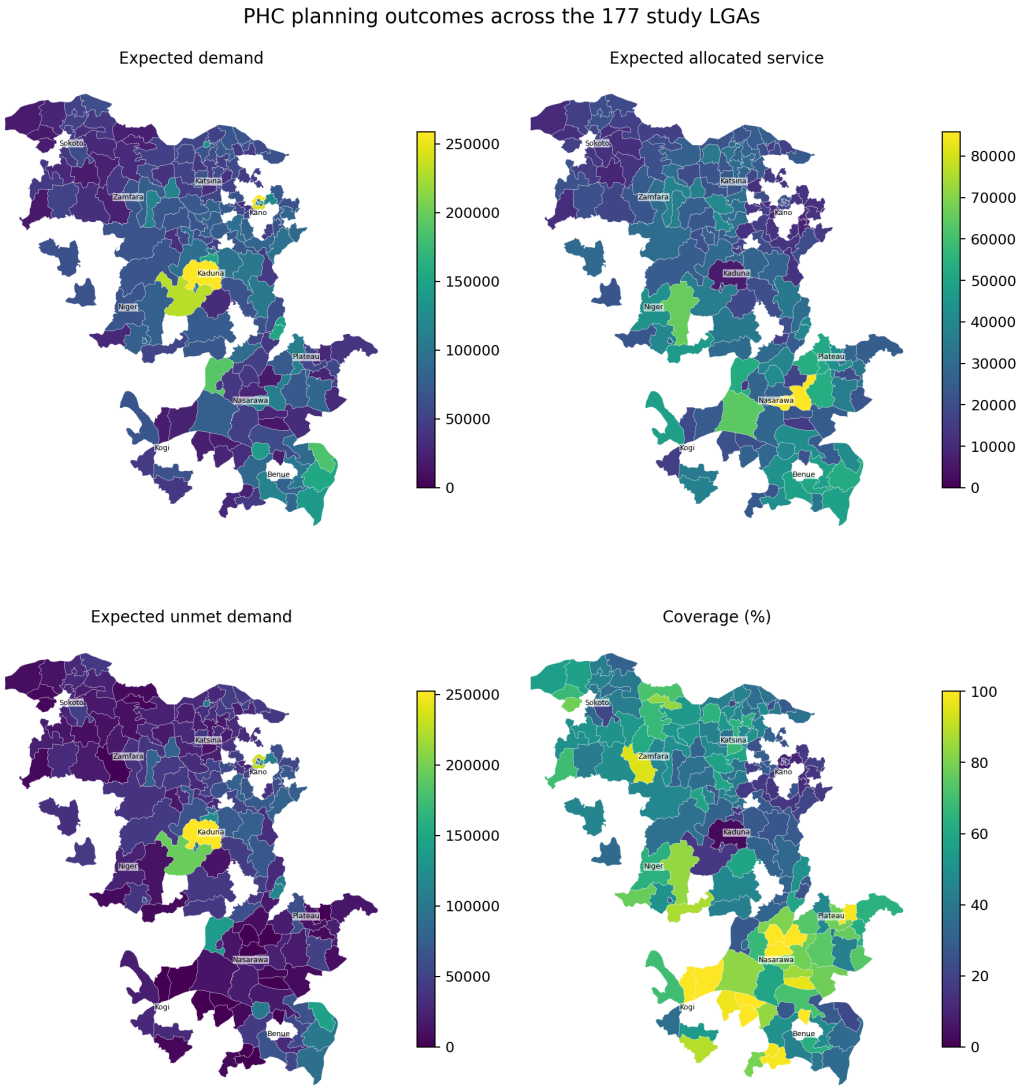


Figure 8: PHC outcomes for the $E00$ solution. Demand (top left), service (top right) and unmet demand (bottom left) are scenario-weighted quarterly visit-equivalents; coverage (bottom right) is their resource-specific ratio at LGA level. Panels use separate value scales, with coverage fixed at 0–100%. White gaps lie outside the common 177-LGA study set.

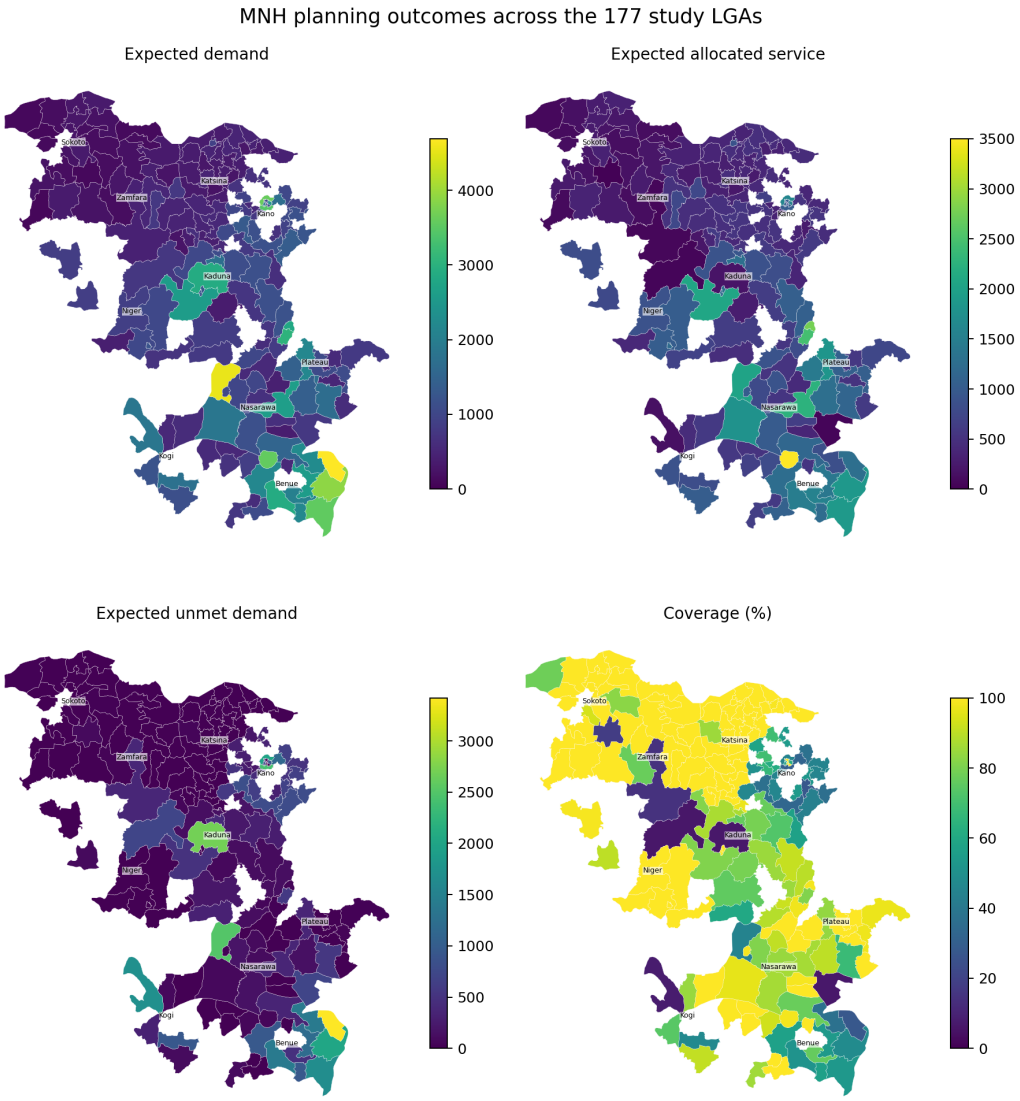


Figure 9: MNH outcomes for the $E00$ solution. Quantities are scenario-weighted quarterly facility-delivery-equivalents. The coverage panel uses the same 0–100% scale as the PHC map, while quantity scales remain resource specific. These are model allocations under the specified network, not observations of patient access

3.10 Held out archetype evaluation of first stage plans

Table 12: Conditional evaluation on the excluded archetype

Held out	**SP loss (million)**	**EV minus SP (cost units)**	**SP minus test-aware (cost units)**	**PHC (%)**	**MNH (%)**
R15	80.958843	129.025	0.081	37.769	71.010
R16	81.085818	125.068	0.079	37.730	70.936
R17	81.019507	126.271	1.966	37.751	70.953
R18	81.110124	125.411	3.156	37.722	70.939

All eight held-out evaluations are feasible. Stochastic plans improve loss by 125.068–129.025 units

over training-mean plans, lying 0.079–3.156 units above test-aware lower bounds (Table 12). Four related folds do not support statistical generalisation.

Later recourse is optimised on the excluded archetype. Shared calibration and nonchronological training limit interpretation to conditional transfer of initial commitments.

3.11 Regional stresses and normative penalty sensitivity

Table 13: Regional and penalty experiments

Case	Own objective (million)	PHC coverage (%)	MNH coverage (%)
Baseline	81.043572	37.743	70.959
Kaduna demand +25%	86.099776	36.246	68.711
Kaduna capacity −50%	82.940178	35.724	76.551
MNH penalty 12.5	80.439492	37.743	70.959
MNH penalty 50	81.903735	37.303	81.443

Kaduna demand stress raises loss by 6.239%, reducing PHC/MNH coverage to 36.246%/68.711%. Halved Kaduna capacity raises loss by 2.340%; PHC falls to 35.724% while MNH rises to 76.551% through reallocation. This does not make disruption beneficial.

Halving the MNH penalty leaves aggregate coverage unchanged at reported precision; doubling it raises MNH to 81.443% while PHC falls to 37.303%. Penalties are normative, not monetary or clinical exchange rates; own-objective values are not directly comparable.

3.12 Discussion

3.12.1 Information value and decision flexibility

Aligned spatial demand, exhausted budgets and shortage-dominated loss accompany small information gains without identifying a single causal bottleneck. Flexibility should be judged against available actions and implementation costs.

3.12.2 Allocation rules and geographic equity

Efficiency and geographic distribution must be reported together. Equal coverage, service floors and vulnerability priorities require separate evaluation. PHC and MNH coverage remain primary because the pooled index combines different service units and cannot establish patient access or fairness.

3.12.3 Resource sensitivity and adverse outcomes

Capacity expansion can favour PHC while eliminating MNH under fixed budgets and stated penalties. Proposition 2 guarantees only expected-loss improvement. Resource-specific safeguards and intervention costs should accompany capacity planning.

3.12.4 Limitations and planning implications

Four related rounds and severity scaling do not span arbitrary relocation. New rounds, dependence ablations and policies for unseen information states are needed. Kaduna stresses are illustrative, not exhaustive.

Quarter-end balances omit clinical deadlines; distance/effective-speed links and readiness-adjusted capacities are proxies. Numerical verification does not validate surveys, population aggregation or

clinical coefficients. Maps show one optimum and VSS one selected mean-demand plan. Prospective planner validation remains necessary.

4 Conclusion

The model separates stochastic initial planning, later adaptation and resource changes across 177 Nigerian LGAs. Information gains are positive but small under the calibrated scenarios. Resource perturbations have larger tested effects; capacity expansion at fixed budgets can eliminate MNH. Held-out results support only conditional transfer of initial commitments.

Exact equal coverage improves local minima while reducing aggregate service and increasing loss. Joint reporting of efficiency, resource-specific service and geographic distribution is essential. Operational conclusions require validated inputs, explicit distributional priorities and evaluation on new spatial conditions.

Acknowledgements

The authors acknowledge the International Organization for Migration Displacement Tracking Matrix and GRID3 for the source data used in this study.

Funding. This research was self-funded by the authors. No grants or external funding were received.

Declarations

Data and code availability: Computational records contain derived inputs, model and analysis scripts, numerical results and calibration checks. Third-party data remain subject to providers' access and redistribution terms. No public repository DOI is asserted.

Use of artificial intelligence: AI was used in code development and preparation of figures. The authors verified all AI-generated materials.

Conflict of interest: The authors declare that they have no financial or non-financial conflicts of interest, financial encumbrances, or obligations that could influence the conduct, findings or publication of this research.

Author contributions: The first author led conceptualisation, methodology, data curation, computation, analysis, visualisation and drafting. The second provided principal guidance, methodology, supervision and critical review. The third contributed to the methodology and reviewed the manuscript. All authors contributed to the review.

References

[1] P.D. Averik, A.A. Obafemi, M. Ogoro, C. Ndabula, G.G. Jidauna, C.B. Akawu, Spatial analysis of physical accessibility to public primary healthcare centers in Kaduna State, Nigeria. FUDMA Journal of Earth and Environmental Sciences, 1(1), 2024, 86–98. https://doi.org/10.33003/jees.2024.0101/07

[2] S. Baas, S. Dijkstra, R.J. Boucherie, A. Zander, A stochastic programming approach for dynamic allocation of bed capacity and assignment of patients to collaborating hospitals during pandemic outbreaks. Health Care Management Science, 29(2), 2026, Article 18. https://doi.org/10.1007/s10729-025-09747-1

[3] S. Bhattarai, Y. Song, Data-driven multi-stage stochastic programming models for integrated hurricane relief logistics and evacuation problem. Computational Management Science, 23(1), 2026, Article 19. https://doi.org/10.1007/s10287-026-00568-3

[4] J.R. Birge, The value of the stochastic solution in stochastic linear programs with fixed recourse. Mathematical Programming, 24(1), 1982, 314–325. https://doi.org/10.1007/BF01585113

[5] J.R. Birge, F. Louveaux, Introduction to stochastic programming (2nd ed.). Springer, 2011. https://doi.org/10.1007/978-1-4614-0237-4

[6] D. Cantor, J. Swartz, B. Roberts, A. Abbara, A. Ager, Z.A. Bhutta, K. Blanchet, D. Madoro Bunte, J.C. Chukwuorji, N. Daoud, W. Ekezie, C. Jimenez-Damary, K. Jobanputra, N. Makhashvili, D. Rayes, M.H. Restrepo-Espinosa, A.J. Rodriguez-Morales, B. Salami, J. Smith, Understanding the health needs of internally displaced persons: A scoping review. Journal of Migration and Health, 4, 2021, Article 100071. https://doi.org/10.1016/j.jmh.2021.100071

[7] Z. Chalabi, D. Epstein, C. McKenna, K. Claxton, Uncertainty and value of information when allocating resources within and between healthcare programmes. European Journal of Operational Research, 191(2), 2008, 530–539. https://doi.org/10.1016/j.ejor.2007.07.026

[8] R.C. Dhungana, S.P. Gupta, I.M. Adhikari, A facility-integrated flow increment problem for emergency evacuation under budget constraints. Operations Research and Decisions, 36(3), 2026, 60–75. https://doi.org/10.37190/ord/218948

[9] W. Ekezie, S. Timmons, P. Myles, P. Siebert, M. Bains, C. Pritchard, An audit of healthcare provision in internally displaced population camps in Nigeria. Journal of Public Health, 41(3), 2019, 583–592. https://doi.org/10.1093/pubmed/fdy141

[10] L.F. Escudero, A. Garín, M. Merino, G. Pérez, The value of the stochastic solution in multistage problems. TOP, 15(1), 2007, 48–64. https://doi.org/10.1007/s11750-007-0005-4

[11] M. Fattahi, E. Keyvanshokooh, D. Kannan, K. Govindan, Resource planning strategies for healthcare systems during a pandemic. European Journal of Operational Research, 304(1), 2023, 192–206. https://doi.org/10.1016/j.ejor.2022.01.023

[12] Federal Ministry of Health and Social Welfare, Nigeria health sector-wide joint annual review 2025: Speech by Muhammad Ali Pate, 12–14 November, 2025. https://health.gov.ng/nigeria-health-sector-wide-joint-annual-review-2025\protect\@normalcr\relax-abuja-nigeria-12-14-november-2025-speech-by-muhammad\protect\@normalcr\relax-ali-pate-the-coordinating-minister-of-health-and-social\protect\@normalcr\relax-welfare-federal-republic-of-nigeria/

[13] Federal Ministry of Health and Social Welfare of Nigeria, National Population Commission, ICF, Nigeria demographic and health survey 2024, 2025. https://health.gov.ng/wp-content/uploads/2025/10/NDHS-2024-Final-Report_Disseminated17102025.pdf

[14] P. Guo, J. Zhu, Coordinating international humanitarian inventory by stochastic dual dynamic programming. Naval Research Logistics, 73(4), 2026, 561–583. https://doi.org/10.1002/nav.70030

[15] H. Heitsch, W. Römisch, Scenario tree reduction for multistage stochastic programs. Computational Management Science, 6(2), 2009, 117–133. https://doi.org/10.1007/s10287-008-0087-y

[16] HiGHS, HiGHS: High-performance parallel linear optimization software, n.d.. https://highs.dev/

[17] I. Kapungu, N.S. Mbare, A.J. Mwakisisile, E.E. Sinkwembe, A stochastic programming approach for asset and liability management of National Health Insurance Fund in Tanzania. Operations Research and Decisions, 36(1), 2026, 74–92. https://doi.org/10.37190/ord/210427

[18] H. Lane, M. Sarkies, J. Martin, T. Haines, Equity in healthcare resource allocation decision making: A systematic review. Social Science & Medicine, 175, 2017, 11–27. https://doi.org/10.1016/j.socscimed.2016.12.012

[19] O. Lawal, F.E. Anyiam, Modelling geographic accessibility to primary health care facilities: Combining open data and geospatial analysis. Geo-spatial Information Science, 22(3), 2019, 174–184. https://doi.org/10.1080/10095020.2019.1645508

[20] C.C. Nnanatu, A. Gadiaga, T.J. Abbott, H. Chamberlain, A.N. Lazar, A.J. Tatem, Modelled gridded population estimates for Nigeria 2025 (Version 3.0) [Data set]. WorldPop, University of Southampton, 2025. https://doi.org/10.5258/SOTON/WP00782

[21] C.G. Rawls, M.A. Turnquist, Pre-positioning of emergency supplies for disaster response. Transportation Research Part B: Methodological, 44(4), 2010, 521–534. https://doi.org/10.1016/j.trb.2009.08.003

[22] R.T. Rockafellar, S. Uryasev, Conditional value-at-risk for general loss distributions. Journal of Banking & Finance, 26(7), 2002, 1443–1471. https://doi.org/10.1016/S0378-4266(02)00271-6

[23] A. Shapiro, D. Dentcheva, A. Ruszczyński, Lectures on stochastic programming: Modeling and theory (3rd ed.). SIAM, 2021. https://doi.org/10.1137/1.9781611976595

[24] A.A. Stinnett, A.D. Paltiel, Mathematical programming for the efficient allocation of health care resources. Journal of Health Economics, 15(5), 1996, 641–653. https://doi.org/10.1016/S0167-6296(96)00493-6

[25] R.M. Stuart, N. Fraser-Hurt, Z. Shubber, L. Vu, N. Cheik, C.C. Kerr, D.P. Wilson, How to do (or not to do)... health resource allocations using constrained mathematical optimization. Health Policy and Planning, 38(1), 2023, 122–128. https://doi.org/10.1093/heapol/czac096

[26] World Bank, Birth rate, crude (per 1,000 people)—Nigeria, 2024 (SP.DYN.CBRT.IN) [Data set]. World Development Indicators, 2026. https://api.worldbank.org/v2/country/NGA/indicator/SP.DYN.CBRT.IN?date=2024&format=json

[27] World Bank, Population, total—Nigeria, 2025 (SP.POP.TOTL) [Data set]. World Development Indicators, 2026. https://api.worldbank.org/v2/country/NGA/indicator/SP.POP.TOTL?date=2025&format=json

[28] X. Yao, K.S. Shehadeh, R. Padman, Multi-resource allocation and care sequence assignment in patient management: A stochastic programming approach. Health Care Management Science, 27(3), 2024, 352–369. https://doi.org/10.1007/s10729-024-09675-6

[29] X. Yin, İ.E. Büyüktahtakın, A multi-stage stochastic programming approach to epidemic resource allocation with equity considerations. Health Care Management Science, 24(3), 2021, 597–622. https://doi.org/10.1007/s10729-021-09559-z

[30] X. Yin, İ.E. Büyüktahtakın, B.P. Patel, COVID-19: Data-driven optimal allocation of ventilator supply under uncertainty and risk. European Journal of Operational Research, 304(1), 2023, 255–275. https://doi.org/10.1016/j.ejor.2021.11.052

[31] C.B. Okechukwu, B.O. Osu, Data-driven spatial demand scenarios for multi-stage stochastic healthcare resource planning under displacement uncertainty. International Journal of Development Mathematics, in press.

Appendix A Demand parameters and source provenance

The PHC anchor divides Q2 2025 BHCPF-facility visits by national population, with different coverage bases [12, 27]. It does not validate LGA utilisation or aggregation. The 2024 annual crude birth rate is 32.493 per 1,000 [26].

NDHS 2024 Table 9.8 reports facility-delivery percentages for live births in the preceding two years [13]. Exact percentages are divided by 100. Observed institutional use can understate unmet obstetric need.

Table A1 Exact state facility-delivery proportions used in primary demand

State	**NDHS source (%)**	**Model proportion**
Benue	59.0	0.590
Kaduna	25.9	0.259
Kano	32.7	0.327
Katsina	15.8	0.158
Kogi	62.2	0.622
Nasarawa	55.7	0.557
Niger	30.2	0.302
Plateau	45.7	0.457
Sokoto	12.5	0.125
Zamfara	15.3	0.153

The 708 primary rows expand to 4,248 terminal records using multipliers 0.90, 1.00 and 1.15. Population is fixed; displacement enters only through the intensity premium.

Excluded subgroup fields are non-comparable; R16 maternal data are unavailable. Source limitations include non-unique site identifiers, category differences, an unresolved R15 date discrepancy and R18 full/common-panel IDP totals of 1,378,124/1,318,021. These concern external validity rather than numerical feasibility.