

Review ARTICLE

Multistage Stochastic Models for Epidemic Resource Allocation: A Comparative Review of Computational Performance and Policy Implications

Shahzaib¹, Muhammad Adnan Munir², Dr. Muhammad Awais Ibrahim Bajwa³, Abdul Wajid⁴, Muhammad Noman⁵

¹Lecturer, Clinical Services, School of Health Sciences, University of Management and Technology, Lahore, Pakistan, Email: shahzaib@umt.edu.pk, ORCID ID: <https://orcid.org/0009-0001-3339-6131>

²Quality Control Analyst, Shaigan Healthcare, Pakistan, Email: adnanmunir201@gmail.com

³Lecturer, Clinical Services, School of Health Sciences, University of Management and Technology, Lahore, Pakistan Email: awais.ibrahim@umt.edu.pk

⁴Lecturer, The University of Haripur, Pakistan, Email: Abdulwajid@uoh.edu.pk

⁵Department of Medical Laboratory Sciences, The University of Faisalabad, Faisalabad, Pakistan. muhammadnoman.mls1@gmail.com, ORCID ID <https://orcid.org/0009-0005-2595-9992>

*Corresponding author: Shahzaib, shahzaib@umt.edu.pk

ARTICLE INFO

Received: 2026-09-15
Revised: 2026-10-08
Accepted: 2026-10-11
Published: 2026-10-13
ID Article: JBSS-2026-0010

ABSTRACT

Background: Epidemic resource allocation requires sequential decisions under uncertain transmission, supply, and behavioral conditions. Published models address these challenges using different objectives and reporting measures, which limits direct comparison. Objective: To organize reported numerical findings from four epidemic resource-allocation model families within a unified mean-risk multistage stochastic programming perspective and to derive transparent comparative indicators without pooling incompatible outcomes. Methods: A structured secondary-data analysis was conducted using reported results from multi-echelon vaccine logistics, equity-aware Ebola allocation, community-level allocation and ventilator planning, and mean-CVaR vaccine-and-treatment logistics. Extracted indicators included predicted infections, infections averted, mixed-integer-programming gap, runtime, stage-wise value of the stochastic solution, mean absolute percentage error, expected deaths, and cost. Runtime speedups and relative cost premium were calculated from reported values. Incomplete observations were not imputed. Results: For the available U.S.-level formulations, predicted infections ranged from 21,515,219 to 21,677,985 and infections averted from 1,702,212 to 1,864,978. Conventional runtimes ranged from 863 to 36,000 s, compared with approximately 12-13 s for decomposition, producing derived speedups of 66-3,000-fold. Stage-wise VSS increased from 0 to 69 objective units. Infection MAPE was 0.06-0.16%, earlier ventilator deployment was associated with approximately 571 fewer expected deaths (1.30%), and an approximately US\$0.15 million cost increase on a US\$70 million baseline represented a derived premium of 0.21%. Conclusion: Decomposition, adaptive stochastic planning, and explicit tail-risk control offer complementary benefits, but the heterogeneous studies cannot be ranked using a single pooled measure. Shared benchmark instances and standardized reporting are needed for confirmatory comparison.

KEYWORDS

Epidemic resource allocation; multistage stochastic programming; CVaR; vaccine logistics; decomposition; VSS; equity



© 2026 The Author(s). Published in the Journal of Biological and Sustainability Sciences (JBSS), a journal of BIOSUSS Publishing (Pvt.) Limited. Authors retain copyright and grant BIOSUSS Publishing (Pvt.) Limited the non-exclusive right of first publication. This open-access article is distributed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). **Journal information and contact:** jbss.biosuss.com

Cite this article: Shahzaib, Munir MA, Bajwa MAI, Wajid A, Noman M. Multistage Stochastic Models for Epidemic Resource Allocation: A Comparative Review of Computational Performance and Policy Implications. J Biol Sustain Sci. 2026;2(1):12-19. DOI: 10.5281/zenodo.23335180

1. Introduction

Epidemic emergencies create a difficult allocation problem because decisions must be made before the scale, location, and duration of transmission are fully known. Vaccines, ventilators, treatment beds, medicines, personnel, and transport capacity are limited, while needs change across regions and over time. A decision that appears efficient at one stage may become inadequate after transmission, supply, or community behavior changes. Consequently, epidemic response increasingly relies on mathematical optimization models that link epidemiological dynamics with resource constraints and permit decisions to adapt as new information becomes available [1-5].

Multistage stochastic programming is especially relevant to this setting. It represents uncertainty through scenarios arranged across successive decision stages and enforces non-anticipativity, meaning that a planner cannot use information before it is observed. The resulting policies can be updated after changes in disease growth or resource availability. The value of the stochastic solution (VSS) compares such a policy with one based on average or expected conditions. A positive VSS indicates that explicitly representing uncertainty improves the modeled objective, although the numerical magnitude remains specific to the objective, disease, and planning horizon [2,6,7].

Recent studies have extended this framework in several directions. Multi-echelon vaccine models integrate national, state, and local distribution decisions and test alternative objective formulations for efficiency and fairness [1]. Equity-aware models incorporate regional disparities when allocating Ebola treatment centers and related resources [2]. Community-level models include behavioral factors such as vaccine hesitancy or examine the timing of ventilator deployment [3,4]. Mean-risk models add Conditional Value-at-Risk (CVaR), allowing planners to balance expected epidemic impact against the consequences of relatively rare but severe scenarios [5,8]. These extensions address different policy questions and therefore should not be interpreted as directly interchangeable experiments.

Beyond vaccine and treatment allocation, epidemic-operations research emphasizes rolling-horizon decisions, spatial targeting, demand-forecast updating, and joint allocation of testing kits, personal protective equipment, and transport capacity. Stochastic and robust approaches are especially valuable when epidemiological parameters, supply availability, and regional demand evolve [9-24]. However, model recommendations can conflict with ethical priorities if total-case minimization overlooks disadvantaged populations. Equity-aware vaccine and protective-equipment studies therefore incorporate coverage, access, or deprivation, while ethical frameworks argue that benefits, harms, and priority to worse-off populations should be evaluated together [2,20,22-25]. These developments reinforce the need for comparisons that preserve both numerical performance and policy meaning.

Computational tractability is central to practical use. Large scenario trees, integer facility decisions, inventory balances, and disease-state transitions can make conventional formulations slow. Decomposition methods reduce this burden by separating a large model into coordinated subproblems. The numerical results considered here show a particularly large runtime difference: alternative formulations required 863 to 36,000 seconds, whereas a decomposition formulation was reported to solve in approximately 12–13 seconds. Such gains are important because outbreak policies may need to be recalculated repeatedly as surveillance and supply data are updated. Nevertheless, solver speed and mixed-integer-programming gap measure optimization performance; they do not by themselves establish epidemiological validity.

Comparison also requires careful treatment of outcome meaning. Predicted infections, infections averted, expected deaths, MAPE, VSS, and monetary cost describe different aspects of model performance. They cannot be combined into a single pooled effect without common populations, horizons, objectives, and uncertainty structures. The present analysis therefore retains each indicator in its reported unit and performs only transparent secondary calculations. Runtime speedup is calculated by dividing the comparator runtime by the reported 12–13-second decomposition range. The relative CVaR cost premium is calculated as the reported US\$0.15 million increment divided by the approximate US\$70 million baseline.

The purpose of this comparative review is to organize the available numerical findings into three complementary tables and two figures. Table 1 reports the available United States supply-chain results; Table 2 presents derived runtime comparisons; and Table 3 summarizes the remaining stochastic, calibration, timing, and risk-sensitive findings. Figure 1 visualizes runtimes on a logarithmic scale, while Figure 2 compares infections averted for the three formulations with complete values in the supplied dataset. Incomplete entries are not estimated. This approach preserves the original meaning of each result, clearly distinguishes reported values from derived calculations, and avoids claiming that the source models were independently rerun.

2. Materials and Methods

2.1 Study design

A structured comparative secondary-data analysis was undertaken to answer three questions: how strongly did solution methods differ in computational performance; what did the reported epidemiological outputs indicate about allocation effectiveness and timing; and how did explicit treatment of uncertainty, equity, and tail risk change the resulting policies? The unit of analysis was a published epidemic resource-allocation model family rather than an individual patient, location, or simulation run. Because the source studies addressed different diseases, populations,

objectives, and planning horizons, the work was designed as a transparent numerical synthesis and not as a conventional meta-analysis.

2.2 Source identification and eligibility

The evidence base consisted of the epidemic-allocation studies and extracted numerical results supplied for this manuscript. Four complementary domains were represented: (i) a multi-echelon vaccine supply-chain model applied to U.S. COVID-19 data [1]; (ii) an equity-aware multistage stochastic model for Ebola treatment-resource allocation [2]; (iii) community-level vaccine-hesitancy and ventilator-allocation models reporting calibration and deployment-timing outcomes [3,4]; and (iv) a mean-CVaR vaccine-and-treatment logistics model [5]. A source was eligible when it reported at least one usable quantitative outcome in computational performance, epidemiological performance, calibration, stochastic value, intervention timing, or risk sensitivity. Duplicate preprint and journal records describing the same model were treated as one model family. Narrative claims without an accompanying numerical value were not used in quantitative comparisons.

2.3 Outcome definitions and extraction

A prespecified extraction structure was used. Computational outcomes comprised runtime in seconds and the reported mixed-integer-programming (MIP) gap. Epidemiological outcomes comprised predicted infections, infections averted, and change in expected deaths. Model-validity outcomes comprised mean absolute percentage error (MAPE). Decision-under-uncertainty outcomes comprised stage-wise value of the stochastic solution (VSS), incremental cost under stronger risk aversion, and the direction of resource reallocation. Values were retained in their original units and linked to their source-model domain so that computational, predictive, clinical, and economic measures were not treated as interchangeable.

Each extracted value was checked for unit consistency and arithmetic compatibility with the associated comparison. Approximate values were preserved as approximate, and reported ranges were retained instead of replaced by point estimates. The supplied record for normalized formulation 2 was incomplete and was excluded rather than reconstructed. No patient-level data, scenario files, random seeds, solver logs, or executable model code were available; consequently, this study did not rerun the optimization models or claim independent computational replication.

2.4 Derived measures

Only calculations directly supported by the extracted numbers were performed. Runtime speedup was defined as comparator runtime divided by decomposition runtime. Because decomposition was reported as 12-13 seconds, speedup was reported as an interval: comparator runtime/13 to comparator runtime/12. The relative CVaR cost premium was defined as incremental cost divided by baseline cost multiplied by 100: (US\$0.15 million / US\$70 million) $\times$ 100 = 0.214%, reported as approximately 0.21%. Absolute differences in infections averted were calculated by subtraction, and relative differences were calculated against the lower comparator value. The midpoint of 12.5 seconds was used solely for plotting the decomposition bar in Figure 1; it was not treated as a newly observed runtime.

2.5 Comparative synthesis

Results were organized into four evidentiary domains: solution quality and runtime; epidemiological output; calibration and timing; and stochastic or risk-sensitive behavior. Within-domain comparisons were descriptive and used absolute differences, percentages, ranges, and ratios. Cross-domain interpretation was narrative. No confidence intervals, hypothesis tests, pooled effect sizes, or formal ranking score were calculated because the underlying populations, objective functions, scenario trees, constraints, and time horizons were not harmonized. Missing observations were shown explicitly and were not imputed.

2.6 Bias control and interpretation rules

Three safeguards were applied to reduce overinterpretation. First, reported values were distinguished from calculations derived in the present analysis. Second, a lower runtime was interpreted as computational efficiency only when considered alongside the reported optimality gap; it was not assumed to imply superior epidemiological performance. Third, calibration error, infections averted, expected deaths, VSS, and monetary cost were interpreted only within their originating model because their denominators and objective scales differed. These rules prevent a cross-study numerical contrast from being mistaken for a head-to-head experiment.

2.7 Ethical considerations

Ethics approval and informed consent were not required because the analysis used published or supplied aggregate numerical results and contained no identifiable participant information. The analysis involved no intervention, recruitment, or access to confidential health records.

3. Results

3.1 U.S.-level formulation performance

Four U.S.-level formulations had at least one usable result, although epidemiological outputs were unavailable for the decomposition formulation (Table 1). Among the three conventional formulations with complete infection outcomes, predicted infections ranged from 21,515,219 to 21,677,985, an absolute spread of 162,766 (0.75% of the highest value). Scalarized formulation 1 produced both the lowest predicted infection count and the highest number

of infections averted (1,864,978). It averted 162,735 more infections than scalarized formulation 2 and 162,766 more than normalized formulation 1, equivalent to approximately 9.56% more than either lower-performing comparator. Scalarized formulation 2 and normalized formulation 1 were epidemiologically almost indistinguishable in the supplied extract: their predicted infection totals differed by 31, and their infections-averted totals also differed by 31. Their optimization performance differed more clearly. Normalized formulation 1 reached a reported 0% MIP gap in 863 seconds, whereas scalarized formulation 2 stopped with a 0.83% gap after 24,000 seconds. Scalarized formulation 1 required 36,000 seconds and reported a 0.14% gap. The decomposition formulation also reported a 0% gap in approximately 12-13 seconds, but the absence of infection outcomes prevents assessment of whether its faster solution preserved the same epidemiological objective value.

Table 1. Reported U.S.-level performance of the available formulations

Method	Predicted infections	MIP gap	Runtime (s)	Infections averted
Scalarized formulation 1	21,515,219	0.14%	36,000	1,864,978
Scalarized formulation 2	21,677,954	0.83%	24,000	1,702,243
Normalized formulation 1	21,677,985	0%	863	1,702,212
Decomposition formulation	Not available in supplied extract	0%	~12–13	Not available in supplied extract

Note: The supplied value for normalized formulation 2 was incomplete and was therefore not estimated or included.

3.2 Computational efficiency

The conventional runtimes spanned 863-36,000 seconds (approximately 14.4 minutes to 10 hours), compared with 12-13 seconds for decomposition. The derived speedup was 66-72-fold relative to normalized formulation 1, 1,846-2,000-fold relative to scalarized formulation 2, and 2,769-3,000-fold relative to scalarized formulation 1 (Table 2). Thus, even the fastest conventional formulation required more than an order of magnitude longer than decomposition, while the slowest required nearly three thousand times longer. These ratios describe computational acceleration within the supplied comparison; they do not establish equivalence of hardware, solver settings, model size, or epidemiological output.

Table 2. Runtime speedups derived from the reported decomposition range

Comparator	Comparator runtime (s)	Speedup using 13 s	Speedup using 12 s
Scalarized formulation 1	36,000	2,769×	3,000×
Scalarized formulation 2	24,000	1,846×	2,000×
Normalized formulation 1	863	66×	72×

3.3 Stochastic value, calibration, timing, and risk sensitivity

The remaining outcomes could not be placed on a common scale and are therefore reported by domain (Table 3). In the equity-aware Ebola model, stage-wise VSS rose from 0 to 69 objective units. Within that model, the result indicates that the advantage of stochastic recourse became larger as additional decisions and uncertainty realizations accumulated; the magnitude should not be interpreted as infections, deaths, or monetary savings. In the community-level analysis, infection MAPE ranged from 0.06% to 0.16%, indicating close agreement for the evaluated observations, although the extract did not provide the sample size or an external validation set.

Timing was consequential in the ventilator-allocation study: earlier deployment was associated with approximately 571 fewer expected deaths, corresponding to a 1.30% reduction in the reported comparison. In the mean-CVaR model, increasing risk aversion raised total modeled cost by approximately US\$0.15 million on a baseline near US\$70 million. This represents a derived premium of 0.21%. The policy composition also changed, with resources shifting from vaccination toward treatment capacity as greater weight was assigned to adverse-tail scenarios. This directional shift is a model result rather than a general recommendation to prioritize treatment over prevention.

Table 3. Cross-study numerical findings retained in their original outcome domains

Model domain	Reported result	Meaning retained in this analysis
Equity-aware Ebola MSSP	Stage-wise VSS increased from 0 to 69 objective units	Modeled value of adaptive stochastic planning increased across the reported stages
Community-level calibration	Infection MAPE: 0.06–0.16%	Reported prediction error for the evaluated observations

Model domain	Reported result	Meaning retained in this analysis
Ventilator deployment timing	~571 fewer expected deaths (1.30%)	Modeled benefit associated with earlier deployment
Mean-CVaR cost sensitivity	~US\$0.15M increment on ~US\$70M baseline	Derived relative cost premium ≈0.21%
Mean-CVaR allocation pattern	Resources shifted from vaccination toward treatment as risk aversion increased	Change in intervention composition under stronger tail-risk protection

3.4 Graphical comparison

Figure 1 displays runtime on a logarithmic scale because the largest reported value was 2,880 times the 12.5-second plotting midpoint. The visual separation is consistent with the speedup intervals in Table 2. Figure 2 shows that scalarized formulation 1 had the highest reported infections-averted value, while scalarized formulation 2 and normalized formulation 1 differed by only 31 infections. No bar was drawn for decomposition because its infections-averted value was unavailable; visual omission therefore represents missing evidence, not zero effect.




Figure 1. Reported runtimes for three conventional formulations and the decomposition formulation. A logarithmic axis is used because values range from approximately 12.5 to 36,000 seconds.




Figure 2. Reported infections averted for the three formulations with complete values in the supplied dataset. No value was imputed for incomplete or unavailable entries.

Abbreviations: CVaR, Conditional Value-at-Risk; MAPE, mean absolute percentage error; MIP, mixed-integer programming; MSSP, multistage stochastic programming; VSS, value of the stochastic solution.

4. Discussion

4.1 Principal findings

This analysis yielded four connected findings. First, decomposition was associated with a 66–3,000-fold reduction in solution time while retaining a reported 0% MIP gap. Second, epidemiological performance did not move in parallel with computational performance: the formulation with the most infections averted was also the slowest, while two substantially faster conventional formulations produced almost identical infection outcomes. Third, increasing stage-wise VSS supported the value of retaining decision flexibility as uncertainty was progressively resolved. Fourth, explicit tail-risk aversion changed both expenditure and allocation composition, demonstrating that risk preference has operational and distributive consequences rather than functioning only as a mathematical tuning parameter.

4.2 Computational implications

The runtime contrast is operationally meaningful. A ten-hour solve may be acceptable for strategic planning but is poorly suited to repeated updates during a rapidly changing outbreak. A 12–13-second solve could support rolling-horizon use, rapid scenario screening, and sensitivity analysis while information is still actionable. The 0% reported MIP gap strengthens the computational result because it suggests that speed was not obtained merely by accepting a visibly weaker bound. Nevertheless, the comparison is not a platform-neutral benchmark. Runtime depends on processor architecture, memory, solver and version, parallelization, presolve, stopping rules, scenario-tree size, and implementation quality. Publication-ready benchmarking should report these items and rerun all methods on the same instances and hardware [7,11,12,14–16].

4.3 Epidemiological and policy interpretation

The infection results expose a practical trade-off between solution speed and the modeled health objective. Scalarized formulation 1 averted about 9.56% more infections than the other complete formulations but required approximately 3.33–9.76 hours longer than the other two conventional formulations. Whether that gain justifies the delay depends on how often the policy must be recomputed, the outbreak growth rate, and the consequences of acting on stale information. Conversely, the 31-infection difference between scalarized formulation 2 and normalized formulation 1 is negligible relative to totals above 1.7 million, whereas their runtimes and reported gaps differ markedly. Decision makers should therefore examine marginal epidemiological gain per unit of computation and not select a method from runtime or infections averted alone.

The ventilator result further shows that allocation is not only a question of quantity and location; timing changes modeled mortality. Earlier deployment may prevent capacity shortages before the epidemic peak, whereas the same units delivered later can have less effect. The rising VSS complements this finding by quantifying the modeled benefit of adapting decisions when uncertainty is revealed. Together, the results favor policies that preserve recourse and schedule repeated re-optimization, while recognizing that their numerical effects remain conditional on transmission parameters, resource lead times, adherence, and the assumed planning horizon [4,6,9–15,17,18,26,27].

4.4 Calibration and validity

The reported infection MAPE of 0.06–0.16% is encouraging for internal predictive agreement, but a small error metric does not by itself demonstrate causal validity or transportability. MAPE can be sensitive to the level and distribution of observed counts, and the supplied extract did not permit examination of residual patterns, out-of-sample performance, parameter uncertainty, or calibration drift. Future studies should pair aggregate error measures with time-resolved residual plots, uncertainty intervals, prospective or external validation, and sensitivity analyses for under-reporting and behavioral change. Optimization results are only as credible as the epidemiological model used to generate their coefficients and state transitions.

4.5 Risk, equity, and portfolio interpretation

The approximately 0.21% CVaR cost premium is small relative to the reported baseline and can be interpreted as the modeled price of greater protection against severe scenarios. Its importance, however, depends on the amount of tail-risk reduction achieved, which was not available on a common quantitative scale in the supplied extract. The accompanying shift from vaccination toward treatment capacity illustrates a portfolio response: when adverse realized outbreaks receive more weight, capacity that mitigates existing severe cases becomes more valuable. This should not be generalized into a rule that treatment should replace prevention. Risk-sensitive allocation should be examined jointly with transmission stage, intervention effectiveness, budget, lead time, and ethical priorities.

Equity remains a separate but interacting concern. An aggregate objective can improve total outcomes while leaving high-risk or poorly served regions behind. Equity constraints may reduce the unconstrained efficiency objective, yet they can improve procedural fairness, geographic coverage, and legitimacy. Future mean-risk analyses should report distributional outcomes by region or population group, show the efficiency-equity frontier, and test whether tail-risk protection benefits the same populations that face the greatest baseline vulnerability [2,5].

Implementation additionally requires transparent governance of the objective function. Public-health agencies should document who selects the risk-aversion parameter, which groups define equity strata, how shortages are

valued, and when a policy may be overridden. Sensitivity analysis should show whether recommended allocations remain stable across plausible transmission, supply, and behavioral assumptions. Participatory review can identify constraints that are mathematically feasible but operationally or ethically unacceptable. These safeguards are important because optimization converts normative choices into numerical weights, and apparently technical settings may redistribute benefits and burdens across regions. Combining model audit trails with explicit ethical criteria would make allocation decisions more accountable, interpretable, and defensible [20,22-26].

4.6 Strengths and limitations

The principal strength is methodological restraint: reported observations, derived calculations, and interpretation are explicitly separated. Missing entries were not imputed; approximate values remained approximate; and heterogeneous outcomes were not collapsed into a misleading composite score. The analysis also translates raw runtimes and costs into decision-relevant ratios while preserving the source units and domains.

Several limitations constrain inference. The source set was purposive and small rather than the product of a systematic search, so selection and publication bias cannot be excluded. The studies differed in disease, geography, intervention, objective function, epidemic model, scenario construction, budget, horizon, and computational environment. Only aggregate reported results were available; source code, instance files, solver logs, repeated-run distributions, and complete uncertainty outputs were absent. Some observations were approximate, and missing infection results for decomposition prevent a direct assessment of its epidemiological equivalence. The calculated speedups therefore describe the supplied study comparison, not a reproducible head-to-head trial. Likewise, VSS units are model-specific, MAPE does not establish external validity, and the 0.21% premium cannot be evaluated fully without a commensurate tail-risk reduction measure.

4.7 Research and implementation priorities

A confirmatory study should implement competing formulations on shared benchmark instances, use identical epidemiological inputs and scenario trees, and run them on the same hardware with prespecified solver tolerances. It should report wall-clock time, optimality gap, objective value, memory use, variability across repeated runs, infections and deaths, resource utilization, equity measures, expected cost, and a common tail-risk metric. Prospective implementation should additionally define update frequency, data latency, governance of risk and equity weights, and decision rules for circumstances in which a faster but slightly less favorable solution is preferable to a slower policy based on outdated information.

5. Conclusions

Epidemic resource allocation benefits from models that can adapt over time, solve quickly, address inequity, and protect against severe scenarios. The available results indicate 66-3,000-fold computational speedups for decomposition, up to 1,864,978 infections averted among the complete *U.S.*-level entries, increasing stage-wise VSS, low reported calibration error, a modeled mortality benefit from earlier ventilator deployment, and an approximately 0.21% cost premium for stronger tail-risk protection. These findings are informative but heterogeneous. They support adaptive, risk-aware allocation while showing that computational efficiency, epidemiological benefit, and fairness must be assessed together. Future work should rerun competing formulations on shared benchmark instances with common epidemiological inputs, scenario trees, hardware, solver settings, equity measures, and reporting standards. Prospective validation, transparent parameter governance, and open code and benchmark data are necessary before these models can be treated as dependable real-time decision-support systems.

Declarations

Author contributions: Shahzaib: Conceptualization, methodology, data curation, and writing—original draft. Abdul Wajid: Conceptualization, literature review, and formal analysis. Muhammad Adnan Munir: Methodology and interpretation of findings. Muhammad Noman: Data validation and interpretation of results. Muhammad Awais Ibrahim Bajwa: Supervision, critical review of the manuscript, and editing. All authors reviewed and approved the final version of the manuscript and agreed to be accountable for all aspects of the work.

Funding: No funding was received for this study from any public, commercial or non-profit organization.

Institutional review board statement: Not applicable, as the study used published and aggregate numerical data only.

Informed consent statement: Not applicable. This study used published aggregate numerical results and involved no human participants or identifiable data.

Data availability: All analyzed numerical data are available in the cited studies and supplied extraction material.

Acknowledgements: None.

Conflicts of interest: The authors declare no conflicts of interest.

Use of artificial intelligence tools: AI tools were used to assist with language editing and manuscript refinement. The authors take full responsibility for the accuracy, originality, and final content of the manuscript.

References

- Jozani K, Sageer NA, Eldardiry H, Tunc S, Buyuktahtakin Toy E. A multi-echelon demand-driven supply chain model for proactive optimal control of epidemics: insights from a COVID-19 study. *arXiv*. 2025;2510.16969. doi: 10.48550/arXiv.2510.16969.
- Yin X, Buyuktahtakin IE. A multi-stage stochastic programming approach to epidemic resource allocation with equity considerations. *Health Care Management Science*. 2021;24:597–622. doi: 10.1007/s10729-021-09559-z.
- Bui H, Eksioğlu S, Proano R, Shen H. Design of targeted community-based resource allocation in the presence of vaccine hesitancy via a data-driven compartmental stochastic optimization model. *arXiv*. 2024;2405.05487. doi: 10.48550/arXiv.2405.05487.
- Yin X, Buyuktahtakin IE, Patel BP. COVID-19: data-driven optimal allocation of ventilator supply under uncertainty and risk. *European Journal of Operational Research*. 2023;304:255–275. doi: 10.1016/j.ejor.2021.11.052.
- Yin X, Buyuktahtakin IE. Risk-averse multi-stage stochastic programming to optimizing vaccine allocation and treatment logistics for effective epidemic response. *IIE Transactions on Healthcare Systems Engineering*. 2022;12:52–74. doi: 10.1080/24725579.2021.1938298.
- Escudero LF, Garin A, Merino M, Perez G. The value of the stochastic solution in multistage problems. *TOP*. 2007;15:48–64. doi: 10.1007/s11750-007-0005-4.
- Birge JR, Louveaux F. *Introduction to Stochastic Programming*. 2nd ed. New York: Springer; 2011. doi:10.1007/978-1-4614-0237-4.
- Rockafellar RT, Uryasev S. Optimization of conditional value-at-risk. *Journal of Risk*. 2000;2:21–41. doi: 10.21314/JOR.2000.038.
- Tanner MW, Sattenspiel L, Ntamo L. Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical Biosciences*. 2008;215:144–151. doi: 10.1016/j.mbs.2008.07.006.
- Yarmand H, Ivy JS, Denton B, Lloyd AL. Optimal two-phase vaccine allocation to geographically different regions under uncertainty. *European Journal of Operational Research*. 2014;233:208–219. doi: 10.1016/j.ejor.2013.08.027.
- Thul L, Powell W. Stochastic optimization for vaccine and testing kit allocation for the COVID-19 pandemic. *European Journal of Operational Research*. 2023;304:325–338. doi: 10.1016/j.ejor.2021.11.007.
- Du M, Sai A, Kong N. A data-driven optimization approach for multi-period resource allocation in cholera outbreak control. *European Journal of Operational Research*. 2021;291:1106–1116. doi: 10.1016/j.ejor.2020.09.052.
- Buyuktahtakin IE, des-Bordes E, Kibis EY. A new epidemics-logistics model: insights into controlling the Ebola virus disease in West Africa. *European Journal of Operational Research*. 2018;265:1046–1063. doi: 10.1016/j.ejor.2017.08.037.
- Long EF, Nohdurft E, Spinler S. Spatial resource allocation for emerging epidemics: a comparison of greedy, myopic, and dynamic policies. *Manufacturing & Service Operations Management*. 2018;20:181–198. doi: 10.1287/msom.2017.0681.
- Liu M, Zhang D. A dynamic logistics model for medical resources allocation in an epidemic control with demand forecast updating. *Journal of the Operational Research Society*. 2016;67:841–852. doi: 10.1057/jors.2015.105.
- Bertsimas D, Boussiou L, Cory-Wright R, et al. From predictions to prescriptions: a data-driven response to COVID-19. *Health Care Management Science*. 2021;24:253–272. doi: 10.1007/s10729-020-09542-0.
- Fadaki M, Abareschi A, Far SM, Lee PTW. Multi-period vaccine allocation model in a pandemic: a case study of COVID-19 in Australia. *Transportation Research Part E*. 2022;161:102689. doi: 10.1016/j.tre.2022.102689.
- Jahani H, Chaleshtori AE, Khaksar SMS, Aghaie A, Sheu JB. COVID-19 vaccine distribution planning using a congested queuing system: a real case from Australia. *Transportation Research Part E*. 2022;163:102749. doi: 10.1016/j.tre.2022.102749.
- Gilani H, Sahebi H. A data-driven robust optimization model by cutting hyperplanes on vaccine access uncertainty in COVID-19 vaccine supply chain. *Omega*. 2022;110:102637. doi: 10.1016/j.omega.2022.102637.
- Wang X, Jiang R, Qi M. A robust optimization problem for drone-based equitable pandemic vaccine distribution with uncertain supply. *Omega*. 2023;119:102872. doi: 10.1016/j.omega.2023.102872.
- Mohammadi M, Dehghan M, Pirayesh A, Dolgui A. Bi-objective optimization of a stochastic resilient vaccine distribution network in the context of the COVID-19 pandemic. *Omega*. 2022;113:102725. doi: 10.1016/j.omega.2022.102725.
- Baloch G, Gzara F, Elhedhli S. COVID-19 PPE distribution planning with demand priorities and supply uncertainties. *Computers & Operations Research*. 2022;146:105913. doi: 10.1016/j.cor.2022.105913.
- Baloch G, Gzara F, Elhedhli S. Risk-based allocation of COVID-19 personal protective equipment under supply shortages. *European Journal of Operational Research*. 2023;310:1085–1100. doi: 10.1016/j.ejor.2023.04.001.
- Donmez Z, Turhan S, Karsu O, Kara BY, Karasan O. Fair allocation of personal protective equipment to health centers during early phases of a pandemic. *Computers & Operations Research*. 2022;141:105690. doi: 10.1016/j.cor.2021.105690.
- Emanuel EJ, Persad G, Kern A, et al. An ethical framework for global vaccine allocation. *Science*. 2020;369:1309–1312. doi: 10.1126/science.abe2803.
- Medlock J, Galvani AP. Optimizing influenza vaccine distribution. *Science*. 2009;325:1705–1708. doi: 10.1126/science.1175570.
- Nguyen C, Carlson JM. Optimizing real-time vaccine allocation in a stochastic SIR model. *PLOS ONE*. 2016;11:e0152950. doi: 10.1371/journal.pone.0152950.