

A Human-Centered AI Coordination Layer for Healthcare Supply Chains During Crises: A Comparative Analysis and Pilot Framework

By Arav Agarwal

AUTHOR BIOGRAPHY

Arav Agarwal is a junior at The Lawrenceville School in Lawrenceville, New Jersey. He is interested in artificial intelligence, public policy, and the healthcare industry and hopes to conduct more projects in these fields in the future.

ABSTRACT

Healthcare resource allocation during crises fails because decisions are made by separate systems operating in silos, limiting the ability to adapt quickly as demand changes. This problem disproportionately affects low- and middle-income countries (LMICs), where a lack of infrastructure leads to more logistical failures. This paper asks how healthcare distributors can improve resource allocation during crises, comparing existing solutions of government stockpiling, mutual aid agreements, logistics management information systems (LMIS), blockchain, localized 3D printing, and a new proposed solution of artificial intelligence (AI) using Ben Shneiderman's reliability, safety, and trust (RST) framework, as well as considerations regarding LMIC feasibility. The paper argues that instead of replacing existing systems, the proposed AI solution should function as a human-monitored coordination layer built on LMIS data, using demand forecasting, shortage prediction, route optimization, and anomaly detection. The paper's original contribution is a system architecture and pilot framework for a regional network of hospitals. This framework includes measurable outcomes such as shortage days, forecast error, response time, emergency shipments, waste, equity gaps, and human overrides. Because no live deployment has been tested, the paper's conclusions are conceptual and not empirical. Yet, the comparisons suggest that a coordinated, transparent, and human-centered AI model may be more realistic and defensible than treating AI as a standalone replacement in an industry where trust is so vital.

Keywords: *healthcare, supply chain, crisis management, artificial intelligence (AI), resource allocation, logistics, transparency, equity, low- and middle-income countries*

INTRODUCTION

1.1 Coordination Failure

During the COVID-19 pandemic, Haiti's healthcare system struggled with shortages of personal protective equipment (PPE), coordination and communication gaps, and a lack of institutional preparedness. These flaws weakened the system's ability to respond to a rapidly changing public health emergency (Faure et al., 2022, p. 1). However, Haiti was not unique, as even highly resourced systems experienced shortages when demand was so erratic (Ranney et al., 2020, paras. 1-2). Therefore, beyond a lack of supplies, crises stem from the problems of knowing what is available, anticipating where demand will rise, moving goods quickly, and fairly deciding which facilities should receive limited resources.

1.2 Scope, Method, and Contribution

This paper discusses the following question: How can healthcare distributors, particularly in LMICs, improve the allocation and distribution of critical medical supplies during crises? By providing a conceptual comparative analysis of six approaches, including five current crisis-response mechanisms and one human-supervised AI model designed to coordinate the complementary functions of each existing response, the paper aims to contribute two distinct elements. First, it frames AI as an integrating and coordinating layer in a wider solution, not a standalone fix. Second, it proposes a specific pilot architecture and evaluation framework to determine whether the model works. The framework also includes specific metrics to determine the AI model's success rate, as the paper does not claim that the proposed AI model will reduce shortages.

The five current approaches were chosen through a narrative literature review of healthcare crisis response and supply-chain operations. They were specifically selected because they represent distinct yet complementary functions that can inform the proposed AI coordination layer.

EXISTING CRISIS-RESPONSE APPROACHES

2.1 Evaluation Framework

An effective solution must be able to function under rapidly changing conditions and prevent patient harm while also providing transparency to hospitals, the government, and the public. For this reason, Shneiderman (2020) proposes principles of reliability, safety, and trust to evaluate AI tools. Reliability concerns whether a solution performs consistently using accurate information. Safety concerns error prevention,

FALL 2026

oversight, and the ability for human intervention when an issue arises. Trust regards transparency, accountability, and whether stakeholders feel that they have the power to question or appeal a decision.

Since these principles were created for AI evaluation, they may not be sufficient for evaluating response mechanisms in LMICs. A system may appear reliable in a controlled setting, but will fail if it is unaffordable, requires specialized staff or digitization, or cannot expand past a pilot program. For this reason, the analysis considers cost, scalability, interoperability, and feasibility.

2.2 Comparative Analysis

Table 1 applies the previous framework to five existing crisis responses. Rather than reviewing each solution in a separate subsection, the table compares the different problems that each approach addresses, its RST evaluation, and practical constraints for implementation in LMICs. Instead of choosing one as a “winner,” the table aims to identify the distinct functions and limitations that an AI coordination layer may address.

Table 1

Comparative Evaluations of Existing Response Mechanisms

Approach	Primary Function	Evaluation Under RST and LMIC Criteria	Role in Coordinated Model
Government stockpiling	Stores critical supplies before a crisis (Bravata et al., 2006, para. 1).	Can provide reliable access when the right products are maintained, but uncertain future demand reduces reliability. Limited cold-storage capacity can make large reserves difficult to sustain in LMICs (Eksioglu et al., 2023, p. 5).	Provides physical inventory for allocation
Mutual aid agreements	Allows hospitals/regions to share resources (den Nijs et al., 2022, para. 1).	Works well when agreements are clear, but reliability fails when all participants face scarcity or organizations prioritize their own reserves. Technology cost is low, but political trust and coordination	Creates pathways for transfers after needs are compared

FALL 2026

		determine feasibility (den Nijs et al., 2022, para. 1).	
LMIS	Tracks inventory, demand, consumption, orders, and shipments across healthcare facilities and distribution centers (Mekonnen et al., 2024, para. 1)	Improves reliability and trust through visibility, but weak cybersecurity or reporting undermines it. More scalable than stockpiling, but necessary technological support causes barriers for LMICs (Kumar et al., 2023, para. 1).	Supplies required data for AI analysis and human oversight
Blockchain	Creates a ledger for transactions and decisions (Kuo et al., 2017, para. 10).	Strengthens integrity and trust by making records for review and auditing, but cannot prove that the original input is accurate or the allocation is fair. Kuo et al. (2017) find that computing infrastructure and maintenance costs limit feasibility.	Functions as an auditing layer for recommendations, approvals, overrides, and transfers
3D printing	Produces selected supplies near the point of need (Mack et al., 2023, para. 16)	Could improve reliability when transportation fails, but safety depends on validated designs, sterilization, and quality control (Mack et al., 2023, para. 1). Equipment, trained staff, regulation, and maintenance restrict scalability in most LMICs.	Provides a contingency production option

Note. The qualitative comparison was constructed by synthesizing the cited literature on five existing healthcare supply chain response mechanisms and evaluating each approach using the RST and LMIC criteria. RST refers to Shneiderman's reliability, safety, and trust framework.

2.3 Comparative Synthesis

The comparison illustrates that no existing approach is a panacea. Stockpiling and 3D printing address the physical availability of products, while mutual aid agreements create a mechanism to share resources. Additionally, LMIS provides stability in inventory and demand, and blockchain strengthens security and data for auditing. These functions operate in different sectors of the problem, so they are not interchangeable.

LMIS is the most important prerequisite for the proposed AI model because accurate coordination depends on accurate information. In settings without digitization, improving inventory reporting, staff training, and basic system connectivity may create more immediate gains than introducing a complex AI or blockchain system that has more possible difficulties.

Therefore, the five existing approaches represent complementary operational capabilities for the AI to coordinate. The AI layer sits above the existing mechanisms and uses shared information to determine which response should be used as conditions change. In this sense, the AI, instead of replacing, coordinates the use of existing tools.

THE PROPOSED AI COORDINATION LAYER

3.1 What the AI System Would Do

The proposed system would combine LMIS and related operational data to forecast demand, predict shortage risk, and recommend allocation and transportation decisions. Data such as recent consumption, patient volume, disease indicators, current stock, open orders, supplier capacity, and lead time would all be input into the AI. Recommendations would consider minimum-stock and equity constraints. Also, the system would flag inconsistent requests or inventory reports, requesting human review.

Research supports the capabilities necessary for this model to function, as AI optimization has shown its ability to prevent hospital drug shortages (Abu Zwaide et al., 2021, pp. 15-17). Moreover, Ahmadi et al. (2022) demonstrate the value of AI inventory methods for perishable pharmaceuticals. Additionally, Dubey et al. (2022) found that AI-driven analytics have shown the ability to strengthen agility and resilience in humanitarian supply chains when supported by data sharing, which AI, acting as a coordination layer for mutual aid agreements, would support.

3.2 Multi-Criteria Decision Logic

Generating an allocation recommendation requires the system to balance priorities that may conflict. To address this complexity, the system could use a fuzzy multi-criteria decision-making approach (Rehman & Ali, 2021, p. 521). A fuzzy reasoning model is useful when categories such as high urgency, access limitations, or supply shortages cannot be accurately represented by a precise cutoff. Instead of looking at each factor in binary terms, the model could assign degrees of priority and produce a ranked allocation recommendation with explanations of the trade-offs.


A fuzzy reasoning model would not determine ethical priorities by itself, as humans would define and approve the criteria, assign weights, create minimum stock floors, and override any issues before deployment. However, reviewers would still be able to alter a recommendation when local knowledge or extenuating circumstances are not reflected in the data provided. While the model would clarify competing priorities, human responsibility would still be vital in deciding how they should be balanced.

3.3 System Architecture

The architecture provided below makes the division of responsibility and system workflow explicit. Data from LMIS and other sources first enter the system, where analytical functions of the AI then forecast demand, predict shortage risk, and detect request anomalies. The outputs of these functions would then be evaluated through fuzzy multi-criteria decision-making (MCDM), considering different factors and tradeoffs. The resulting recommendation would then be sent to a human reviewer, who could approve, modify, or reject the final allocation decision. Decisions, overrides, and transfers would all be recorded for auditing.

The AI architecture separates the different analytical functions from the decision process. This way, a singular model is not forced to perform the entire process, allowing for better tailoring to specific tasks. Since fuzzy MCDM uses human-defined criteria, human reviewers remain responsible for establishing the criteria, reviewing recommendations, addressing cases that fall outside the scope of the AI, and approving important decisions. This modular structure allows AI to coordinate different functions while maintaining human accountability.

Figure 1
Human-Centered AI Coordination Process



Note. This conceptual workflow was made from the proposed AI architecture described in Section 3.3. The figure organizes the system into data inputs, separate AI analytical functions, fuzzy MCDM-based recommendations, human review, implementation, and auditing over the entire system.

3.4 RST and Feasibility

Reliability depends on data quality, whether that is inventory reports, product codes, or consumption data, as these all can produce false confidence. The system should therefore display data-quality warnings and confidence intervals. Also, it should avoid making precise recommendations when the information provided is weak or inconsistent.

Safety requires layered control, as the model is a form of decision support, not autonomous work behind the scenes. Human reviewers must be able to inspect and understand the reasons for an allocation recommendation, compare alternatives, and override the result when needed. A manual fallback process should

also remain readily available for a possible connectivity failure or unexpected AI behavior. Last, high-impact allocation changes should require additional review.

Beyond technical transparency, for trust, stakeholders must know who is responsible, how to appeal a decision, and how errors are corrected. Successful AI implementation depends heavily on stakeholder concerns and perceived accountability (Hogg et al., 2023, para. 5). In the proposed model, trust would be built through published decision criteria, explanations, overrides, and independent audits, leading to performance reports.

PROPOSED PILOT AND EVALUATION FRAMEWORK

4.1 Pilot Setting

First, a practical pilot should be regional rather than national, including one distribution hub and a network that already uses a basic LMIS. The facilities should vary in size and geography so that the pilot can test urban and rural transportation differences, as well as ensure fair allocation. Before activating the AI model, the network should collect baseline data for later comparison.

4.2 Operating Model

During the pilot, the system would generate daily recommendations and provide alerts when a certain supply falls below the minimum stock level. Any change in decision would require a short reason code to help train the AI and provide more data. The system would record whether the recommendation was followed and what happened afterward. This structure is consistent with human-centered AI principles and stakeholder-centered implementation (Hogg et al., 2023, para. 4; Shneiderman, 2022).

The pilot should be phased. Phase one would run the model in shadow mode; Phase two would allow recommendations for a limited group of low-risk products; and Phase three would expand to more critical items if the necessary accuracy, safety, and governance standards are reached. This phased sequence prevents the use of an untested model in critical situations and on critical supplies (Vasey et al., 2022, para. 4).

4.3 Quantitative Evaluation

The pilot will be evaluated against the baseline period without AI and against other networks to find gaps in the coordination. As Abu Zwaida et al. (2021) and Ahmadi et al. (2022) describe, inventory studies provide useful precedent for evaluation, shortage prevention, forecasting, and inventory performance. The

FALL 2026

following measures make the proposal testable and define what would need to be improved before the model could be scaled.

Table 2
Proposed Pilot Performance Measures

Metric	Definition	Purpose
Shortage days	Number of days a critical item is unavailable, reported per 1,000 facility days	Tests whether availability improves
Forecast error	Difference between predicted and actual consumption at a facility or throughout the network	Tests prediction and analytical reliability
Response time	Time from a shortage report to an approved allocation or transfer	Tests crisis response speed
Emergency shipments	Number and cost of unplanned urgent deliveries	Tests whether planning improves
Waste and expiration (Ahmadi et al., 2022)	Quantity and value of products discarded	Tests inventory efficiency
Equity Gap	Difference in service level or shortage days between urban and rural facilities	Tests distributional fairness
Override rate	Percentage of AI recommendations changed by human reviewers, with reason codes attached	Tests model issues
False-positive alerts	Percentage of anomaly flags to be unsupported after human review	Tests whether monitoring for anomalies creates an unnecessary burden
System uptime and	Availability of the platform and percentage of required	Tests operational

FALL 2026

data completion	data fields reported on time	feasibility
-----------------	------------------------------	-------------

Note. These proposed performance measures were constructed by translating the evaluation objectives described in Section 3.4 into measurable indicators, as well as using the waste and expiration metrics that Ahmadi et al. (2022) discuss. Therefore, the pilot could be evaluated against a pre-AI baseline across availability, reliability, response speed, efficiency, equity, human oversight, and feasibility. Quantitative outcomes should be supplemented by interviews/surveys to assess whether AI explanations are understandable, whether there is an added burden, user trust, and overall coordination (Hogg et al., 2023, para. 2). Results should be separated by facility type and geography, so average improvement does not hide unequal outcomes.

DISCUSSION

5.1 Infrastructure

Since the AI model serves as a coordination layer, the proposed system is only as strong as the operational environment around it. Therefore, investments in LMIS, connectivity, staff training, and basic warehouse needs may produce more value than investment in a complex model (Kafoe, 2024, para. 3-4). AI should only be added after the network can support it.


However, this does not mean that LMICs must wait for perfect infrastructure. Beyond the AI being a phased model, it can operate with uncertainty and use conservative recommendations when data is weak. Instead of trying to optimize the entire supply chain, the AI could prioritize a small number of critical supplies, especially important during crises.

5.2 Stakeholder Cooperation and Conflict

Healthcare workers, hospital administrators, distributors, suppliers, governments, regulators, and patients all have legitimate but different priorities. The governance structure should therefore specify and publicize which objectives are prioritized, who may change these priorities, and how a facility can appeal an allocation decision. While the AI can clarify trade-offs in a decision, it should not choose the moral rule. Humans decide how these constraints are defined as a prerequisite for model operation.

Figure 2

Position of the Proposed System in a Human-Centered AI Framework



Note. The figure was made by adapting Shneiderman's human-centered AI framework to illustrate the proposed AI coordination layer according to its relative levels of computer automation and human control. The red marker represents the proposed system. Adapted from *Human-Centered AI* (p. 60), by B. Shneiderman, 2022, Oxford University Press. Copyright 2022 by Ben Shneiderman.

5.3 Risks and Mitigation

Bias is a serious risk because historical healthcare data reflect unequal access (Celi et al., 2022, para. 5). A facility serving an underserved population may report lower consumption simply because patients could not afford care, perpetuating the cycle. Furthermore, Obermeyer et al. (2019) demonstrate how a health-management algorithm produced racial bias when cost was used as a metric for medical need. Therefore, proper data, equity constraints, subgroup testing, minimum stock protections, and regular audits are a necessity.

Hospitals or suppliers may also attempt to manipulate the system, but unusual requests should be compared with patient volume, past consumption, and neighboring facilities. However, an anomaly should never remove a hospital from consideration for allocation, given the erratic nature of crises. Instead, there should be a documented verification process and a human decision, utilizing blockchain to make changes visible.

Technical failures must be treated as expected, as cyberattacks, bad data, software defects, and network outages may occur. The pilot should maintain a manual allocation process, back up data, and a clear authority to suspend the model in this situation.

FALL 2026

LIMITATIONS AND FUTURE RESEARCH

First, this paper only addresses the conceptual design of an AI system. The pilot architecture and metrics provided make the proposal testable, but they do not demonstrate that the model will reduce shortages. Furthermore, the paper does not provide a full cost model, and cost and staffing requirements may differ greatly across regions.

Future research should begin with past inventory and crisis data, followed by first-phase deployment in shadow mode and a regional pilot. Researchers should compare the proposed system with current allocation rules, examine performance under missing data, and test whether equity constraints improve underserved facilities without causing major damage elsewhere. Additional work should also study whether hospital staff understand explanations, when reviewers override recommendations, and whether transparency improves trust or causes gridlock. Since trust is so essential to the success of such a program, studying user behavior is vital.

CONCLUSION

Healthcare supply chain crises cannot be solved by one technology, as every response offers something different. Therefore, AI's most defensible role is coordination. Built on LMIS data, an AI system can forecast demand, predict shortages, optimize transportation, and flag unusual requests. However, AI should not have the capability to make high-stakes decisions without review, so humans must define the values, resolve stakeholder conflicts, and remain accountable for outcomes.

The main contribution of this paper is a human-supervised architecture and a measurable pilot plan. The proposed system has five stages: data input, AI analysis, human review, final allocation, and audit. In the analysis stage, fuzzy model reasoning would allow for competing goals to be evaluated and trade-offs to be made instead of only valuing a single objective. The system's success would be judged through shortage days, forecast accuracy, response time, emergency shipments, waste, equity gaps, override patterns, false positives, and operational reliability.

ACKNOWLEDGEMENTS

I would like to thank Professor Amol Joshi of Wake Forest University for his mentorship and guidance throughout the development of this paper.

I also acknowledge the use of ChatGPT to generate Figure 1. I provided the figure's content, structure, and intended relationships and asked ChatGPT to convert that information into a visual diagram. I reviewed and approved the final figure, while preserving my own ideas and information.

REFERENCES

- Abu Zwaïda, T., Pham, C., & Beauregard, Y. (2021). Optimization of inventory management to prevent drug shortages in the hospital supply chain. *Applied Sciences*, 11(6), 2726. <https://doi.org/10.3390/app11062726>
- Ahmadi, E., Mosadegh, H., Maihami, R., Ghalekhondabi, I., Sun, M., & Süer, G. A. (2022). Intelligent inventory management approaches for perishable pharmaceutical products in a healthcare supply chain. *Computers & Operations Research*, 147, 105968. <https://doi.org/10.1016/j.cor.2022.105968>
- Bravata, D. M., Zaric, G. S., Holty, J.-E. C., Brandeau, M. L., Wilhelm, E. R., McDonald, K. M., & Owens, D. K. (2006). Reducing mortality from anthrax bioterrorism: Strategies for stockpiling and dispensing medical and pharmaceutical supplies. *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science*, 4(3), 244-262. <https://doi.org/10.1089/bsp.2006.4.244>
- Celi, L. A., Cellini, J., Charpignon, M.-L., Dee, E. C., Dérnoncourt, F., Eber, R., Mitchell, W. G., Moukheiber, L., Schirmer, J., Situ, J., Paguio, J., Park, J., Wawira, J. G., & Yao, S. (2022). Sources of bias in artificial intelligence that perpetuate healthcare disparities—A global review. *PLOS Digital Health*, 1(3), e0000022. <https://doi.org/10.1371/journal.pdig.0000022>
- den Nijs, K., Edivaldo, J., Châtel, B., Uleman, J., Olde Rikkert, M., Wertheim, H., & Quax, R. (2022). A global sharing mechanism of resources: Modeling a crucial step in the fight against pandemics. *International Journal of Environmental Research and Public Health*, 19(10), 5930. <https://doi.org/10.3390/ijerph19105930>
- Dubey, R., Bryde, D. J., Dwivedi, Y. K., Graham, G., & Foropon, C. (2022). Impact of artificial intelligence-driven big data analytics culture on agility and resilience in humanitarian supply chain: A practice-based view. *International Journal of Production Economics*, 250, 108618. <https://doi.org/10.1016/j.ijpe.2022.108618>
- Eksioglu, S. D., Proano, R. A., Kolter, M., & Nurre Pinkley, S. (2023). Designing drone delivery networks for vaccine supply chain: A case study of Niger. *IIEE Transactions on Healthcare Systems Engineering*, 14(3), 193-213. <https://doi.org/10.1080/24725579.2023.2268113>
- Faure, J. A., Wang, C.-W., Chen, C.-H. S., & Chan, C.-C. (2022). Assessment of the functional capacity and preparedness of the haitian healthcare system to fight against the covid-19 pandemic: A narrative review. *Healthcare*, 10(8), 1428. <https://doi.org/10.3390/healthcare10081428>
- Hogg, H. D. J., Al-Zubaidy, M., Talks, J., Denniston, A. K., Kelly, C. J., Malawana, J., Papoutsis, C., Teare, M. D., Keane, P. A., Beyer, F. R., & Maniatopoulos, G. (2023). Stakeholder perspectives of clinical artificial intelligence implementation: Systematic review of qualitative evidence. *Journal of Medical Internet Research*, 25, e39742. <https://doi.org/10.2196/39742>
- Kafoe, A. S. (2024). Supply chain resilience strategy for healthcare organizations. *Advances in Logistics, Operations, and Management Science*, 1-54. <https://doi.org/10.4018/979-8-3693-4288-6.ch001>

-
- Kumar, A., Mani, V., Jain, V., Gupta, H., & Venkatesh, V. (2023). Managing healthcare supply chain through artificial intelligence (AI): A study of critical success factors. *Computers & Industrial Engineering*, 175, 108815. <https://doi.org/10.1016/j.cie.2022.108815>
- Kuo, T.-T., Kim, H.-E., & Ohno-Machado, L. (2017). Blockchain distributed ledger technologies for biomedical and health care applications. *Journal of the American Medical Informatics Association*, 24(6), 1211-1220. <https://doi.org/10.1093/jamia/ocx068>
- Mack, K., Hofmann, M., Lakshmi, U., Cao, J., Auradkar, N., Arriaga, R., Hudson, S., & Mankoff, J. (2023). Rapid convergence: The outcomes of making PPE during a healthcare crisis. *ACM Transactions on Computer-Human Interaction*, 30(1), 1-25. <https://doi.org/10.1145/3542923>
- Mekonnen, B. A., Worku, M. C., & Tefera, B. B. (2024). Evaluation of logistics management information system and availability of non-program tracer drugs in public health facilities in bahir dar city, northwest ethiopia. *PLOS ONE*, 19(4), e0302319. <https://doi.org/10.1371/journal.pone.0302319>
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464), 447-453. <https://doi.org/10.1126/science.aax2342>
- Ranney, M. L., Griffeth, V., & Jha, A. K. (2020). Critical supply shortages — the need for ventilators and personal protective equipment during the covid-19 pandemic. *New England Journal of Medicine*, 382(18). <https://doi.org/10.1056/nejmp2006141>
- Rehman, O. U., & Ali, Y. (2021). Enhancing healthcare supply chain resilience: Decision-making in a fuzzy environment. *The International Journal of Logistics Management*, 33(2), 520-546. <https://doi.org/10.1108/ijlm-01-2021-0004>
- Shneiderman, B. (2020). Human-Centered artificial intelligence: Reliable, safe & Trustworthy. *International Journal of Human-Computer Interaction*, 36(6), 495-504. <https://doi.org/10.1080/10447318.2020.1741118>
- Shneiderman, B. (2022). *Human-centered AI*. Oxford University Press.
- Vasey, B., Nagendran, M., Campbell, B., Clifton, D. A., Collins, G. S., Denaxas, S., Denniston, A. K., Faes, L., Geerts, B., Ibrahim, M., Liu, X., Mateen, B. A., Mathur, P., McCradden, M. D., Morgan, L., Ordish, J., Rogers, C., Saria, S., Ting, D. S. W., . . . McCulloch, P. (2022). Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: DECIDE-AI. *BMJ*, 377, e070904. <https://doi.org/10.1136/bmj-2022-070904>